

# ***RTView® Enterprise User's Guide***

Version 5.1.1



RTView®

Copyright © 1998-2019. All rights reserved.

No part of this manual may be reproduced, in any form or by any means, without written permission from Sherrill-Lubinski Corporation. All trademarks and registered trademarks mentioned in this document are property of their respective companies.

RTView Enterprise®

©2013-2019 Sherrill-Lubinski Corporation. All Rights Reserved.

No part of this manual may be reproduced, in any form or by any means, without written permission from Sherrill-Lubinski Corporation. All trademarks and registered trademarks mentioned in this document are property of their respective companies.

#### LIMITATIONS ON USE

Use, duplication, or disclosure by the U.S. Government is subject to restrictions as set forth in the Technical Data - Commercial Items clause at DFARS 252.227-7015, the Rights in Data - General clause at FAR 52.227-14, and any other applicable provisions of the DFARS, FAR, or the NASA FAR supplement.

SL, SL-GMS, GMS, RTView, SL Corporation, and the SL logo are trademarks or registered trademarks of Sherrill-Lubinski Corporation in the United States and other countries. Copyright © 1998-2019 Sherrill-Lubinski Corporation. All Rights Reserved.

JMS, JMX and Java are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries. They are mentioned in this document for identification purposes only.

No part of this manual may be reproduced, in any form or by any means, without written permission from Sherrill-Lubinski Corporation.

All trademarks and registered trademarks mentioned in this document are property of their respective companies.



SL Corporation  
240 Tamal Vista Blvd.  
Corte Madera, CA 94925 USA

Phone: 415.927.8400  
Fax: 415.927.8401  
Web: <http://www.sl.com>

<b>Preface</b> .....	1
About This Guide .....	1
Audience .....	1
Document Conventions .....	1
Additional Resources .....	2
Release Notes .....	2
Support Knowledge Base .....	2
SL Documentation .....	2
Documentation for IBM and RTView Enterprise .....	2
Documentation for Infrastructure and RTView Enterprise .....	3
Documentation for Kafka and RTView Enterprise.....	4
Documentation for Oracle and RTView Enterprise.....	4
Documentation for Solace and RTView Enterprise.....	5
Documentation for TIBCO and RTView Enterprise.....	5
Contacting SL.....	6
Internet .....	6
Technical Support.....	7
<b>Chapter 1 - Introduction to RTView Enterprise</b> .....	9
Displays .....	10
RTViewCentral .....	10
RTView DataServers.....	11
Solution Packages .....	12
System Requirements .....	13
<b>Chapter 2 - Using the Monitor</b> .....	15
Login to RTView Enterprise.....	15
Display Server Login .....	15
<b>HTML Login</b> .....	16
Overview .....	17
Navigation Tabs .....	18
SERVICE TREE Tab .....	19
SERVICE VIEWS Tab.....	24
COMPONENTS Tab.....	24
ALERTS Tab.....	27
ADMIN Tab.....	30
CUSTOM Tab .....	31
Fundamental Structure of Displays .....	31
Heatmaps .....	32
Mouse-over .....	34
History Heatmaps .....	34
Tables .....	35
Multiple Column Sorting .....	35
Column Visibility.....	36

Column Filtering .....	36
Column Locking .....	38
Column Reordering .....	38
Saving Settings .....	38
Row Paging .....	39
Row Color Code .....	39
Row Keyboard Selection .....	40
Trend Graphs .....	40
Mouse-over .....	41
Log Scale .....	41
Time Range .....	41
Popup Menu .....	42
Export PDF Report .....	42
Title Bar .....	43
Multiple Windows .....	44
RTViewCentral Displays .....	45
All Management Areas .....	45
Area Heatmap .....	46
Area Table .....	47
Multi Area Service Views .....	49
Group/Service Heatmap .....	50
Group/Region Heatmap .....	52
Group / Service Table .....	53
Services CI Type Summary .....	55
Services History Heatmap .....	58
Single Area Service Views .....	60
Single Area: Group/Service Heatmap .....	61
Single Area: Region/Service Heatmap .....	62
Single Area: Group / Service Table .....	64
Single Area: Services CI Type Summary .....	66
Single Area: Services History Heatmap .....	69
Service Summary Views .....	71
Service By CI Type .....	71
Service Summary .....	75
Service Health Heatmap .....	79
Key Metrics Views .....	81
Service KM Heatmap .....	82
Service KM Table .....	86
Service KM History .....	89
Service KM History (Alt) .....	93
Available KM Metrics and Alerts .....	98
Component Views .....	109
CI / Service Tree View .....	109
CI / Service Table .....	111
CI / Type Heatmap .....	113
CI / Type Table .....	115

Metric Explorer .....	117
Metric Explorer.....	117
Alert Views .....	122
RTView Alerts Table .....	122
Alerts Table - HTML .....	126
Alert History Table.....	128
Administration .....	131
Alert Administration .....	131
Tabular Alert Administration.....	134
Alert Administration - HTML .....	137
Alert Overrides Administration - HTML.....	139
Alert Admin Audit .....	140
Alert Action Audit Trail .....	142
CMDB Administration.....	143
CMDB Admin .....	143
CMDB Administration - HTML .....	149
Add CIs to the CMDB .....	150
Delete a CI.....	150
Edit CI Properties .....	151
<b>Create, Delete, Rename or Merge Owners, Areas, Groups or Services</b> .....	151
<b>Move a Service, Group or Area</b> .....	151
Architecture .....	154
System Overview .....	154
RTView Data Servers .....	156
Data Server Summary .....	158
RTView History Table Statistics.....	160
RTView Cache Tables .....	161
RTView CI Stats Tables .....	163
RTView CI Type Defs.....	164
RTView KM Defs .....	166
About .....	167
Property Views .....	168
Properties Configuration .....	168
System Properties .....	170
Applied Properties .....	171
All Properties .....	174
Properties Descriptions.....	176
Diagram Views .....	178
Steps to Create a Diagram Display.....	180
Create an Object Template Display .....	181
Node Administration Display .....	182
Link Admin Display .....	184
Diagram Properties Admin Display .....	185
Add Diagrams to your Project.....	187
View Diagram Displays.....	187
Optional Diagram Display Customizations .....	187

Other - HTML Displays.....	189
Common/Alerts - HTML.....	189
Alerts Table by Component - HTML .....	189
Alert Detail for Component - HTML.....	190
Alert Configuration for Component - HTML .....	192
Common/System - HTML .....	194
System Overview - HTML.....	194
RTView Data Servers - HTML.....	195
Data Server Summary - HTML.....	196
CI Type Definitions - HTML .....	197
CI Stats - HTML .....	198
Cache Table - HTML.....	199
<b>Chapter 3 - RTView DataServer for IBM .....</b>	<b>203</b>
IBM DB2.....	204
DB2 Instances.....	204
All Instances Heatmap .....	205
All Instances Table .....	209
Instance Summary .....	210
DB2 Members.....	213
All Members Table .....	213
All Members Grid.....	215
Member Summary .....	216
All DB2 Databases .....	219
All Databases Table .....	219
All Databases Grid .....	221
Single Database .....	223
Database Summary .....	223
Partition Table .....	226
Partition Heatmap .....	228
Partition Grid .....	231
Reports.....	233
Trends .....	234
DB2 Log .....	236
IBM DB2 - HTML.....	236
IBM DB2 Overview Display - HTML.....	237
DB2 Instances View - HTML.....	238
DB2 Instances Table - HTML .....	238
DB2 Instances Heatmap - HTML .....	238
Instance Summary - HTML .....	239
DB2 Members View - HTML .....	242
DB2 Members Table - HTML.....	242
DB2 Member Summary - HTML .....	244
DB2 Databases View - HTML.....	246
DB2 Databases Table - HTML .....	246
Single DB2 Database Summary - HTML .....	249

Partition Table - HTML.....	252
Partition Heatmap - HTML.....	254
IBM MQ .....	258
Overview .....	258
IBM MQ Brokers View .....	260
All Brokers Grid .....	260
All Broker Connections .....	262
Single Broker Summary .....	263
Channels View .....	266
All Channels Table .....	266
Single Channel Summary .....	268
Single Channel Detail .....	272
Queue Managers.....	274
All Queue Managers Table .....	274
Queues.....	275
All Queues Table .....	275
Single Queue Summary.....	277
IBM MQ - HTML .....	280
IBM MQ Overview Display - HTML.....	281
IBM MQ Brokers View - HTML .....	282
IBM MQ Broker Connections Table - HTML.....	282
IBM MQ Brokers Heatmap - HTML.....	282
IBM MQ Broker Summary - HTML .....	283
IBM MQ Channels View - HTML .....	285
IBM MQ Channels Table - HTML.....	285
IBM MQ Channels Heatmap - HTML.....	286
IBM MQ Channel Summary - HTML .....	287
IBM MQ Queues View - HTML .....	289
IBM MQ Queue Managers Table - HTML .....	289
IBM MQ Queues Table - HTML .....	290
IBM MQ Queues Heatmap - HTML.....	291
IBM MQ Queue Summary - HTML.....	292
IBM WebSphere.....	294
Overview .....	294
Key Features .....	295
Metrics for WebSphere Server .....	295
End-to-End Context for WebSphere.....	295
All WebSphere Servers View .....	296
All Servers Heatmap .....	296
All Servers Table .....	298
All Servers Grid .....	300
All Servers Detail.....	301
Single Server View .....	305
Server Summary .....	306
JDBC Providers .....	309
Thread Pool .....	312

All Applications View .....	313
All Sessions Heatmap.....	314
Session Charts By App .....	315
Session Detail By App .....	317
All Applications Detail.....	319
Single Application View.....	320
Application Summary .....	321
Component Detail.....	323
Module Totals - Charts .....	325
Module Totals - Tables .....	326
IBM WebSphere - HTML.....	328
WebSphere Overview - HTML.....	328
WebSphere Servers Heatmap - HTML.....	330
WebSphere Server Summary - HTML .....	330
WebSphere Apps Table - HTML.....	331
WebSphere Apps Heatmap - HTML.....	332
WebSphere App Summary - HTML .....	332
<b>Chapter 4 - RTView DataServer for Infrastructure .....</b>	<b>333</b>
Amazon Web Services .....	334
EC2 Instances .....	335
Amazon EC2 Instance Heatmap.....	335
Amazon EC2 Instance Table.....	337
Amazon EC2 Instance Summary .....	339
Amazon Web Services - HTML.....	341
Amazon EC2 Overview - HTML .....	342
Amazon EC2 Instance Heatmap - HTML .....	343
Amazon EC2 Instance Table - HTML .....	345
Amazon EC2 Instance Summary - HTML .....	347
Docker .....	348
Engine View .....	349
Engines Heatmap .....	349
Engines Table .....	352
Engine Summary.....	354
Container View .....	356
Containers Heatmap .....	356
Containers Table .....	359
Container Summary.....	361
Docker - HTML .....	364
Docker Overview - HTML.....	365
Docker Engines View - HTML .....	366
Docker Engines Table - HTML.....	367
Docker Engines Heatmap - HTML.....	367
Docker Single Engine Summary - HTML.....	370
Docker Containers View - HTML .....	371

	Docker Containers Table - HTML.....	371
	Docker Containers Heatmap - HTML.....	372
	Docker Single Container Summary- HTML.....	374
JBoss .....		375
	JBoss Servers View .....	375
	All Servers Heatmap .....	376
	All Servers Table .....	377
	Server Summary .....	379
	JBoss Applications.....	381
	Applications Heatmap .....	381
	Applications Summary.....	384
JBoss - HTML .....		386
	JBoss Overview - HTML .....	387
JBoss Servers - HTML.....		387
	JBoss Servers Table - HTML.....	387
	JBoss Servers Heatmap - HTML.....	387
	JBoss Single Server Summary - HTML.....	388
JBoss Applications - HTML.....		388
	JBoss Apps Table - HTML.....	388
	JBoss Apps Heatmap - HTML.....	388
	JBoss App Totals - HTML .....	389
MongoDB.....		389
	Mongo Instance View .....	389
	All Instances Heatmap .....	390
	All Instances Table .....	392
	Single Instance Summary.....	395
	Mongo Database View .....	398
	All Databases Heatmap .....	398
	All Databases Table .....	401
	Database Summary .....	403
	Mongo Collection View.....	405
	All Collections Heatmap.....	405
	All Collections Table.....	408
	Collection Summary.....	410
MongoDB - HTML.....		412
	MongoDB Overview - HTML .....	413
	Mongo Instances View - HTML .....	414
	MongoDB Instances Table - HTML.....	415
	MongoDB Instances Heatmap - HTML.....	415
	MongoDB Instance Summary - HTML .....	418
	Mongo Databases View - HTML .....	419
	MongoDB Databases Table - HTML.....	419
	MongoDB Databases Heatmap - HTML.....	420
	MongoDB Single Database Summary - HTML.....	422
	Mongo Collections View - HTML.....	423
	MongoDB Collections Table - HTML .....	423

MongoDB Collections Heatmap - HTML .....	424
MongoDB Single Collection Summary - HTML .....	426
.....	427
MySQL Database .....	427
All MySQL Databases.....	427
All Servers Heatmap .....	427
All Servers Table .....	430
Single MySQL Database .....	432
Server Summary .....	433
Servers Properties .....	435
Servers Operations .....	436
User Tables .....	438
MySQL Database - HTML.....	439
MySQL Overview Display- HTML .....	439
All MySQL Instances View - HTML .....	440
Instances Heatmap- HTML .....	440
All MySQL Instances - HTML.....	443
Single MySQL Instance View - HTML.....	443
Single MySQL Instance Summary- HTML .....	444
Instance Properties- HTML .....	445
Instance Operations- HTML.....	445
Instance User Tables- HTML.....	446
MS SQL Server .....	447
All Servers View.....	448
All Servers.....	448
All Servers Heatmap .....	450
Single SQL Server View .....	452
Server Summary .....	453
Database Details .....	457
Wait Stats .....	459
Performance Trends.....	461
Database/Table Sizes.....	464
MS SQL - HTML .....	466
MS SQL Overview - HTML.....	466
MS Servers View - HTML .....	467
MS SQL Servers Table - HTML .....	468
MS SQL Servers Heatmap - HTML .....	468
MS SQL Server View - HTML.....	470
MS SQL Server Summary - HTML .....	471
MS SQL Server Performance - HTML .....	473
MS SQL Server Waits Table - HTML.....	474
MS SQL Server DB Table Sizes - HTML .....	475
Node.js .....	475
Node/Master View.....	476
Node Master Table.....	476

Node Master Summary .....	479
Node Request View .....	481
Node Requests Table .....	481
Node Request Summary .....	483
Process View .....	485
All (Node) Processes Heatmap.....	486
All (Node) Processes Table.....	488
(Node) Process Summary .....	490
Node.js - HTML.....	492
Node.js Overview .....	493
Node/Master View - HTML .....	494
Node/Master Info Table - HTML .....	495
Node Master Summary - HTML .....	496
Node Requests View - HTML .....	497
Node Requests Table - HTML.....	497
Node Request Summary - HTML .....	498
Process View - HTML .....	499
Node Processes Table - HTML.....	499
Node Processes Heatmap - HTML.....	500
Node Process Summary - HTML.....	501
RTView Host Agent.....	503
All Hosts .....	503
All Hosts Heatmap .....	503
All Hosts Table .....	505
All Hosts Grid.....	508
All Processes Table .....	509
All Network Table .....	511
All Storage Table .....	513
Single Host .....	514
Host Summary .....	515
RTView Host Agent - HTML.....	517
Hosts Overview - HTML .....	518
Hosts Heatmap - HTML.....	519
Hosts Table - HTML.....	520
Host Summary - HTML .....	521
Host Processes - HTML .....	522
Host Network Interfaces - HTML .....	523
Host Storages - HTML .....	524
VMware vCenter .....	525
Clusters View .....	526
All Clusters .....	526
Hosts View .....	529
All Hosts .....	529
Single Host Summary .....	532
Host Health .....	535
Host NICs.....	537

Virtual Machines View .....	539
All VMs Heatmap .....	539
All VMs Table .....	542
All VMs Disk Table .....	545
Single VM Summary .....	547
Datstores View .....	550
All Datstores Table .....	551
Hosts by Datastore Table .....	553
VMs by Datastore Table .....	554
Single Datastore Summary .....	556
Networks View .....	558
All Networks Table .....	558
Events/Alarms View .....	560
All Events Table .....	561
All Alarms Table .....	563
VMware vCenter - HTML .....	564
VMware Overview Display .....	565
VMware Hosts View - HTML .....	566
VMware Hosts Table - HTML .....	567
VMware Host Summary - HTML .....	568
VMware Machines View - HTML .....	570
VMware Virtual Machines Table - HTML .....	570
VMware Virtual Machine Summary - HTML .....	571
<b>Chapter 5 - RTView DataServer for Kafka .....</b>	<b>573</b>
Kafka .....	573
Kafka Clusters View .....	574
All Clusters Table .....	575
Cluster Performance .....	577
Kafka Topics View .....	579
All Topics Heatmap .....	579
All Topics Table .....	581
Single Topic Summary .....	584
All Topics for Cluster .....	588
Single Topic for Cluster .....	590
Kafka Brokers View .....	592
All Brokers Heatmap .....	592
All Brokers Table .....	595
All Broker Metrics Table .....	598
Single Broker Summary .....	606
Kafka Zookeepers View .....	615
All Zookeepers Heatmap .....	615
All Zookeepers Table .....	618
Zookeepers Summary .....	620
Kafka Producers View .....	625
All Producers Heatmap .....	625

All Producers Table .....	628
Producer Summary .....	632
Kafka Consumers View .....	637
All Consumers Heatmap .....	637
All Consumers Table .....	639
Consumers Summary .....	642
Kafka - HTML .....	647
Kafka Overview - HTML.....	648
Kafka Clusters View - HTML.....	649
Kafka Clusters Table - HTML .....	650
Kafka Single Cluster Performance - HTML .....	650
Kafka Brokers View - HTML .....	651
Kafka Brokers Table - HTML.....	652
Kafka Brokers Heatmap - HTML.....	652
Kafka Single Broker Summary - HTML.....	655
Kafka Single Broker JVM Runtime Summary - HTML.....	656
Kafka Single Broker Topics Summary - HTML .....	657
Kafka Single Broker Topic Lag Summary - HTML.....	658
Kafka Zookeepers View - HTML.....	658
Kafka Zookeepers Table - HTML .....	659
Kafka Zookeepers Heatmap - HTML .....	659
Kafka Single Zookeeper Summary - HTML .....	662
Kafka Single Zookeeper JVM Runtime Summary - HTML.....	664
Kafka Topics View - HTML.....	666
Kafka Topics Table - HTML.....	666
Kafka Topics Heatmap - HTML.....	667
Kafka Single Topic Summary - HTML.....	670
Kafka Single Topic Partition Summary - HTML .....	672
Kafka Topics Activity by Cluster - HTML .....	672
Kafka Brokers Activity by Topic - HTML.....	673
Kafka Producers View - HTML .....	674
Kafka Producers Table - HTML.....	675
Kafka Producers Heatmap - HTML.....	675
Kafka Single Producer Summary - HTML.....	678
Kafka Single Producer JVM Runtime Summary - HTML .....	680
Kafka Consumers View - HTML.....	682
Kafka Consumers Table - HTML .....	682
Kafka Consumers Heatmap - HTML .....	683
Kafka Single Consumer Summary - HTML .....	686
Kafka Single Consumer JVM Runtime Summary - HTML .....	688
Kafka Single Consumer Lag Summary - HTML .....	690
<b>Chapter 6 - RTView DataServer for Oracle .....</b>	<b>691</b>
Oracle Coherence.....	692
Cluster Selector .....	693
Cluster Views .....	694

Cluster - Overview .....	695
Caches / Nodes / Alerts .....	699
Memory/Network Health .....	702
Stability Metrics .....	704
All Services History .....	706
All Caches History .....	709
All Nodes History .....	713
Proxy Services .....	715
Proxy / Extend Overview .....	716
Proxy / Extend Connections .....	717
Proxy / Extend Detail .....	723
Proxy Nodes History .....	727
Extend Connections History .....	729
Cache Services .....	730
Single Service Summary .....	731
Service Metrics Overview .....	735
Service Metric Heatmap .....	736
Single Service History .....	738
Cache Service Detail .....	740
Federated Clusters .....	742
Federated Destination Detail .....	743
Federated Destination Summary .....	746
Federated Origin Detail .....	748
Federated Origin Summary .....	751
All Caches .....	753
All Caches Heatmap .....	754
Storage Nodes Cache Map .....	755
Current Size Chart .....	757
Current Activity Chart .....	758
Single Cache .....	760
Single Cache Summary .....	760
Size Trends .....	763
Activity Trends .....	766
Cache Detail Tables .....	768
Storage Manager Detail .....	770
Node/Group Distribution .....	772
Front/Back Analysis .....	774
All Nodes .....	775
All Nodes by Type/Host/Memory .....	776
All Nodes CPU .....	777
All Nodes Grid View .....	778
Communication Issues .....	779
All Nodes - Detail .....	781
Invocation Service Detail .....	783
Single Node .....	785
Node Summary .....	785
Service Trends .....	788

Node Detail .....	790
JVM Summary.....	792
JVM Memory Trends.....	794
JVM GC Trends.....	796
System Properties .....	798
Time Range Analysis .....	799
Service Comparison .....	799
Cache Comparison.....	800
OC Administration.....	802
OC Metrics Administration.....	802
Cluster MBean Servers .....	804
Management Settings .....	806
Node Administration .....	808
Cache Administration .....	810
RTView Cache Tables .....	811
About .....	813
.....	813
Oracle Coherence - HTML.....	813
Oracle Database Overview Display- HTML .....	814
Oracle Databases View - HTML.....	815
Oracle Databases Table - HTML .....	815
Databases Heatmap - HTML.....	816
Databases Summary - HTML .....	819
Oracle Instances View - HTML.....	820
Oracle Instances Table - HTML .....	820
Instances Heatmap - HTML.....	821
Instance Summary - HTML .....	823
Oracle Database .....	824
Database Instances View .....	824
All Databases Heatmap .....	824
All Databases Table .....	826
Database Summary .....	828
All Instances Heatmap .....	830
All Instances Table .....	833
Instance Summary .....	836
Database Details View .....	838
Space .....	838
Contention .....	841
Processes.....	843
Sessions .....	845
Queries.....	847
Reports.....	850
Oracle Database - HTML .....	852
Oracle Database Overview Display- HTML .....	852
Oracle Databases View - HTML.....	854
Oracle Databases Table - HTML .....	854

Databases Heatmap - HTML.....	855
Databases Summary - HTML.....	858
Oracle Instances View - HTML.....	859
Oracle Instances Table - HTML.....	859
Instances Heatmap - HTML.....	860
Instance Summary - HTML.....	862
Connector for Oracle Enterprise Manager .....	863
Oracle WebLogic.....	863
WebLogic Servers View.....	864
All Servers Heatmap.....	864
All Servers Table.....	867
All Servers Grid.....	870
All Clusters Table.....	872
Single WebLogic Server View.....	873
Server Summary.....	874
WLS JVM Summary.....	877
WLS Server Detail.....	880
WLS JDBC Summary.....	886
WLS ThreadPool Summary.....	890
Work Manager.....	893
Persistent Stores.....	895
Application Views View.....	897
Cluster Apps Table.....	897
Cluster App Summary.....	899
Server Apps Heatmap.....	902
Server Apps Summary.....	905
Server Apps Trends.....	908
App Components Heatmap.....	910
WebLogic Single Application Summary.....	912
App Components Summary.....	916
JMS Servers View.....	918
JMS Servers Grid.....	919
JMS Server Summary.....	920
JMS Metric Trends.....	924
JMS Server Detail.....	926
JMS Destinations View.....	929
Totals by Destination.....	929
Detail by Destination.....	933
Distribution by Server.....	936
Destination by Server.....	939
Detail by Server.....	942
JMS Bridges View.....	944
All Bridges Table.....	945
Bridge Summary.....	947
JMS Connections View.....	949
JMS Connections - Detail by Server.....	950

Oracle WebLogic - HTML .....	953
WebLogic Overview Display - HTML .....	953
WebLogic Servers View - HTML .....	956
WebLogic Servers Table - HTML.....	956
WebLogic Servers Heatmap - HTML.....	957
WebLogic Clusters Table - HTML .....	959
Single WebLogic Server View - HTML .....	959
WebLogic Server Summary - HTML .....	960
WebLogic JVM Summary - HTML.....	961
WebLogic Server Detail Tables - HTML.....	962
WebLogic JDBC Summary - HTML .....	963
WebLogic Server Work Manager Table - HTML .....	965
WebLogic JMS Persistent Stores Detail Tables - HTML.....	966
WebLogic Applications View - HTML .....	967
WebLogic Cluster Applications Table - HTML.....	968
WebLogic Clustered Application Summary - HTML.....	969
WebLogic Server Applications Heatmap - HTML .....	970
WebLogic Server Applications Summary - HTML.....	972
WebLogic Application Summary - HTML .....	973
WebLogic Application Metric Trends - HTML.....	975
WebLogic Application Components Heatmap - HTML.....	976
WebLogic Application Component Summary - HTML .....	978
.....	979
<b>Chapter 7 - RTView DataServer for Solace .....</b>	<b>981</b>
Brokers .....	982
Brokers Overview .....	982
Brokers Heatmap .....	984
Brokers Table .....	986
Broker Summary .....	994
Broker Sensors .....	996
Broker Provisioning.....	997
Broker Interface.....	999
Brokers Message Spool .....	1001
CSPF Neighbors .....	1003
Neighbors Table .....	1004
Neighbors Diagram .....	1005
Neighbors Summary .....	1006
VPNs.....	1008
VPNs Heatmap .....	1008
VPNs Table .....	1012
VPNs Summary .....	1016
Clients.....	1019
Clients Table.....	1020
Client Summary .....	1025
Bridges.....	1028

Bridges Table.....	1028
Bridges - Diagram .....	1032
Bridge Summary .....	1033
Endpoints .....	1036
Endpoints Table .....	1036
Endpoint Summary .....	1038
Capacity .....	1041
Capacity Table .....	1041
Capacity - Summary .....	1044
Capacity Trends .....	1045
Syslog Events.....	1046
Syslog Events Table.....	1047
.....	1049
<b>Chapter 8 - RTView DataServer for TIBCO.....</b>	<b>1051</b>
TIBCO ActiveMatrix .....	1052
TIBCO ActiveSpaces .....	1052
All Realms Views.....	1052
All Realms Heatmap.....	1052
All Realms Table.....	1054
Single Realm Views .....	1056
Realm Summary .....	1056
Realm Servers .....	1059
All Node Views.....	1061
All Nodes Heatmap .....	1061
All Nodes Table .....	1063
Single Node Views .....	1067
Node Summary .....	1068
All Proxy Views.....	1071
All Proxies Heatmap.....	1071
All Proxies Table.....	1073
Single Proxy Views.....	1076
Proxy Summary .....	1076
All Keeper Views .....	1080
All Keepers Heatmap.....	1080
All Keepers Table.....	1082
Single Keeper Views .....	1083
Keeper Summary .....	1084
TIBCO ActiveSpaces - HTML .....	1086
Realms Views - HTML .....	1088
TIBCO ActiveSpaces Realms Table - HTML .....	1088
TIBCO ActiveSpaces Realms Heatmap - HTML .....	1089
TIBCO ActiveSpaces Realm Summary - HTML.....	1090
TIBCO ActiveSpaces Realm Servers - HTML.....	1092
Nodes Views - HTML.....	1093
TIBCO ActiveSpaces Nodes Table - HTML.....	1093

TIBCO ActiveSpaces Nodes Heatmap - HTML.....	1094
TIBCO ActiveSpaces Node Summary - HTML .....	1095
Proxies Views - HTML .....	1096
TIBCO ActiveSpaces Proxies Table - HTML .....	1097
TIBCO ActiveSpaces Proxies Heatmap - HTML .....	1097
TIBCO ActiveSpaces Proxy Summary - HTML.....	1099
Keepers Views - HTML .....	1100
TIBCO ActiveSpaces StateKeepers Table - HTML.....	1100
TIBCO ActiveSpaces StateKeepers Heatmap - HTML.....	1101
TIBCO ActiveSpaces Keeper Summary - HTML.....	1102
TIBCO ActiveSpaces (2.x) .....	1103
Spaces View.....	1103
All Metaspaces Table.....	1103
Metaspace Summary.....	1106
All Spaces Table .....	1109
All Spaces Heatmap .....	1112
Space Summary .....	1114
All Queries Table .....	1118
Query Summary.....	1120
Members View .....	1122
All Members Table .....	1123
All Members Heatmap .....	1126
Member Summary .....	1129
Member Summary - Process .....	1133
Member Summary -JVM .....	1136
Members by Space Table .....	1139
Members by Space Heatmap .....	1142
Spaces by Member Table .....	1144
Member by Space Summary .....	1146
TIBCO ActiveSpaces (2.x) - HTML .....	1150
TIBCO ActiveSpaces 2 Overview .....	1150
MetaSpaces - HTML.....	1152
TIBCO ActiveSpaces 2 Metaspaces Table - HTML.....	1152
TIBCO ActiveSpaces 2 MetaSpace Summary - HTML.....	1153
Spaces - HTML .....	1154
TIBCO ActiveSpaces 2 Spaces Table - HTML.....	1154
TIBCO ActiveSpaces 2 Spaces Heatmap - HTML.....	1155
TIBCO ActiveSpaces 2 Space Summary - HTML .....	1156
Members - HTML .....	1158
TIBCO ActiveSpaces 2 All Members Table - HTML.....	1158
TIBCO ActiveSpaces 2 Members Heatmap - HTML.....	1159
TIBCO ActiveSpaces 2 Member Summary - HTML.....	1161
TIBCO Adapters.....	1163
All Adapters View.....	1163
All Adapters Heatmap .....	1163
All Adapters Table .....	1166

Single Adapter View .....	1168
Adapter Summary .....	1169
TIBCO Adapters - HTML .....	1172
TIBCO Adapters Overview - HTML .....	1172
All Adapters View - HTML .....	1173
TIBCO Adapters Table - HTML .....	1174
TIBCO Adapters Heatmap - HTML .....	1174
TIBCO Adapter Summary - HTML.....	1176
TIBCO BusinessEvents.....	1178
Clusters / Nodes View.....	1178
Clusters .....	1178
Cluster Summary .....	1182
Cluster Nodes Table.....	1185
Cluster Nodes Heatmap.....	1188
Inference Node Summary .....	1191
Storage Node Summary .....	1196
Events / Concepts View .....	1200
Agent Events .....	1200
Agent Event Summary .....	1203
Event Cache Hits.....	1206
Event Hit Summary .....	1208
Concept Cache Hits.....	1212
Concept Hit Summary .....	1213
Channels.....	1216
All Inference Agents.....	1218
All RTC Reports .....	1221
TIBCO BusinessEvents - HTML .....	1223
TIBCO BusinessEvents Overview - HTML.....	1224
BE Clusters - HTML .....	1226
TIBCO BE Clusters Table - HTML.....	1226
TIBCO BE Clusters Heatmap - HTML.....	1227
TIBCO BE Cluster Summary - HTML .....	1229
BE Nodes - HTML .....	1230
TIBCO BE Cluster Nodes Table - HTML.....	1230
TIBCO BE Cluster Nodes Heatmap - HTML.....	1231
TIBCO BE Inference Node Summary - HTML.....	1232
TIBCO BE Storage Node Summary - HTML.....	1234
BE Events - HTML .....	1236
TIBCO BE Events Table - HTML.....	1236
TIBCO BE Agent Event Summary - HTML .....	1237
TIBCO BE Event Cache Hits Table - HTML .....	1239
TIBCO BE Event Hit Summary - HTML .....	1240
BE Concepts - HTML.....	1242
TIBCO BE Concepts Table - HTML .....	1242
TIBCO BE Concept Hit Summary - HTML.....	1243
TIBCO BE Channel Status Table - HTML.....	1244

TIBCO BE Inference Agents Table - HTML .....	1245
TIBCO BE RTC Txn Manager Reports - HTML .....	1246
.....	1247
TIBCO BusinessWorks.....	1247
BW Applications.....	1248
BW All Applications Heatmap.....	1248
BW All Applications Table.....	1251
BW Single Application Summary .....	1254
BW Containers.....	1257
All Containers Heatmap .....	1257
All Containers Table .....	1260
Single Container Summary .....	1263
BW AppNodes .....	1265
BW All AppNodes Heatmap .....	1265
BW All AppNodes Table .....	1268
BW Single AppNode Summary.....	1270
BW AppSlices .....	1273
BW All AppSlices Heatmap .....	1273
BW All AppSlices Table .....	1276
BW Single AppSlice Summary .....	1278
BW Processes .....	1281
BW All Processes Heatmap.....	1281
BW All Processes Table.....	1284
BW Single Process Summary .....	1287
BW5 Engines.....	1290
All Engines Heatmap.....	1290
All Engines Table .....	1293
All Engines Grid.....	1296
Single Engine Summary .....	1298
BW5 Processes .....	1301
All Processes Heatmap .....	1301
All Processes Table .....	1305
Single Process Summary .....	1308
BW5 Activities .....	1310
All Activities Heatmap .....	1310
All Activities Table .....	1313
Single Activity Summary .....	1316
BW5 Servers .....	1318
All Servers Heatmap .....	1318
All Servers Table .....	1321
All Servers Grid .....	1322
Single Server Summary.....	1324
Server Processes.....	1326
Single Server Process - Summary .....	1327
TIBCO BusinessWorks - HTML.....	1329
TIBCO BusinessWorks Overview Display - HTML.....	1329

BW Applications - HTML .....	1331
TIBCO BusinessWorks Applications Table - HTML .....	1331
TIBCO BusinessWorks Applications Heatmap - HTML .....	1332
TIBCO BusinessWorks Application Summary - HTML.....	1334
BW Containers - HTML.....	1336
TIBCO BusinessWorks Containers Table - HTML .....	1336
TIBCO BusinessWorks Containers Heatmap - HTML .....	1337
TIBCO BusinessWorks Container Summary - HTML.....	1340
BW Application Nodes - HTML.....	1341
TIBCO BusinessWorks Application Nodes Table - HTML .....	1341
TIBCO BusinessWorks Application Nodes Heatmap - HTML .....	1342
TIBCO BusinessWorks Application Node Summary - HTML.....	1343
BW Application Slices - HTML.....	1345
TIBCO BusinessWorks Application Slices Table - HTML.....	1345
TIBCO BusinessWorks Application Slices Heatmap - HTML.....	1346
TIBCO BusinessWorks Application Slice Summary - HTML .....	1348
BW Processes - HTML .....	1351
TIBCO BusinessWorks Processes Table - HTML .....	1351
TIBCO BusinessWorks Processes Heatmap - HTML .....	1352
TIBCO BusinessWorks Process Summary - HTML .....	1354
BW Activities - HTML .....	1356
TIBCO BusinessWorks Activities Table - HTML .....	1356
TIBCO BusinessWorks Activities Heatmap - HTML .....	1357
TIBCO BusinessWorks Activity Summary - HTML.....	1358
.....	1360
TIBCO BusinessWorks 5 Monitor - HTML.....	1360
TIBCO BusinessWorks 5 Overview Display .....	1360
BW5 Servers - HTML .....	1362
TIBCO BusinessWorks 5 Servers Table - HTML.....	1362
TIBCO BusinessWorks 5 Servers Heatmap - HTML.....	1363
TIBCO BusinessWorks 5 Server Summary - HTML.....	1365
BW5 Engines - HTML .....	1366
TIBCO BusinessWorks 5 Engines Table - HTML .....	1366
TIBCO BusinessWorks 5 Engines Heatmap - HTML .....	1367
TIBCO BusinessWorks 5 Engine Summary - HTML.....	1370
BW5 Processes - HTML .....	1372
TIBCO BusinessWorks 5 Processes Table - HTML.....	1372
TIBCO BusinessWorks 5 Processes Heatmap - HTML.....	1373
TIBCO BusinessWorks 5 Process Summary - HTML.....	1376
BW5 Activities - HTML .....	1377
TIBCO BusinessWorks 5 Activities Table - HTML.....	1378
TIBCO BusinessWorks 5 Activities Heatmap - HTML.....	1379
TIBCO BusinessWorks 5 Activity Summary - HTML.....	1381
.....	1383
TIBCO Enterprise Message Service .....	1383
All EMS Servers .....	1383

All Servers Heatmap .....	1383
All Servers Table .....	1386
All Servers Grid .....	1390
All Servers Topology .....	1392
Single EMS Server .....	1396
Single Server Summary.....	1396
Single Server Trends.....	1401
Single Server Tables .....	1402
EMS Topics.....	1407
All Topics Heatmap .....	1407
All Topics Table .....	1412
All Topics Summary .....	1415
Single Topic Summary .....	1418
Single EMS Topic-Clients .....	1422
Single Topic By Server .....	1425
EMS Queues.....	1427
All Queues Heatmap .....	1427
All Queues Table .....	1432
All Queues Summary .....	1435
Single Queue Summary .....	1439
Single EMS Queue-Clients.....	1443
Single Queue By Server.....	1446
EMS Clients.....	1448
Connections.....	1449
Bridges, Users, Ports.....	1451
Routes .....	1453
Producers.....	1457
Producer Summary .....	1459
Consumers .....	1462
Consumer Summary .....	1465
Durables .....	1468
TIBCO Enterprise Message Service - HTML .....	1470
<b>TIBCO EMS Overview - HTML</b> .....	1471
EMS Servers - HTML.....	1473
TIBCO EMS Servers Table - HTML .....	1473
TIBCO EMS Servers Heatmap - HTML .....	1474
TIBCO EMS Server Summary - HTML .....	1477
TIBCO EMS Server Trends - HTML .....	1479
TIBCO EMS Server Destinations - HTML .....	1480
TIBCO EMS Server Connections - HTML .....	1481
TIBCO EMS Bridges, Users, Ports - HTML .....	1482
EMS Topics - HTML.....	1483
TIBCO EMS Topics Table - HTML .....	1484
TIBCO EMS Topics for Server Summary - HTML .....	1485
TIBCO EMS Topics Heatmap - HTML .....	1486
TIBCO EMS Topic Summary - HTML .....	1489

TIBCO EMS Topic Detail by Server - HTML .....	1490
EMS Queues - HTML.....	1491
TIBCO EMS Queues Table - HTML .....	1492
TIBCO EMS Queues for Server Summary - HTML .....	1492
TIBCO EMS Queues Heatmap - HTML.....	1493
TIBCO EMS Queue Summary - HTML.....	1496
TIBCO EMS Queue Detail By Server - HTML.....	1497
EMS Routes - HTML.....	1497
TIBCO EMS Routes - HTML.....	1498
TIBCO EMS Route - HTML.....	1498
EMS Producers - HTML.....	1499
TIBCO EMS Producers - HTML .....	1500
TIBCO EMS Producer - HTML.....	1501
EMS Consumers - HTML.....	1501
TIBCO EMS Consumers - HTML.....	1502
TIBCO EMS Consumer - HTML.....	1502
EMS Durables - HTML .....	1504
TIBCO EMS Durables - HTML.....	1504
TIBCO EMS Durable - HTML .....	1505
.....	1506
TIBCO FTL .....	1506
FTL Servers.....	1506
All Servers Heatmap .....	1506
All Servers Table .....	1509
All Group Servers Table.....	1513
All Satellites Table .....	1514
Server Summary .....	1517
FTL Clients.....	1520
All Clients Heatmap .....	1521
All Clients Table .....	1523
Clients by Group .....	1526
Client Summary .....	1528
Client Metrics.....	1531
FTL Events.....	1532
Events .....	1532
TIBCO FTL - HTML .....	1535
TIBCO FTL Overview - HTML .....	1535
FTL Servers - HTML.....	1536
FTL Servers Heatmap - HTML.....	1537
FTL Servers Table - HTML.....	1539
FTL Server Summary - HTML .....	1540
FTL Clients - HTML.....	1541
All Clients Heatmap - HTML.....	1541
All Clients Table - HTML.....	1543
Single Client Summary - HTML.....	1544
TIBCO Hawk .....	1546

Hawk Agents Table .....	1546
Hawk Alerts Table.....	1547
RTView Agent Administration .....	1549
TIBCO Hawk - HTML.....	1550
Hawk Overview - HTML .....	1550
Hawk Agents Table - HTML .....	1551
Hawk Alerts Table - HTML.....	1552
<b>Chapter 9 - RTView Manager.....</b>	<b>1555</b>
Login to RTView Manager.....	1555
Overview .....	1556
RTView Servers Displays .....	1556
Data Servers .....	1556
Display Servers.....	1559
Historian Servers.....	1560
Version Info.....	1562
RTView Servers Displays - HTML .....	1564
Data Servers - HTML.....	1564
Data Server Summary - HTML.....	1565
Display Servers - HTML .....	1567
Display Server Summary - HTML .....	1568
Historian Servers - HTML.....	1568
JVM Displays.....	1569
All JVMs.....	1570
All JVMs Heatmap.....	1570
All JVMs Table.....	1571
Single JVM .....	1574
JVM Summary.....	1574
JVM System Properties.....	1578
JVM Memory Pool Trends .....	1579
JVM GC Trends.....	1583
JVM Displays - HTML .....	1585
JVM Overview - HTML .....	1585
.....	1586
JVMs Table - HTML .....	1586
.....	1587
JVMs Heatmap - HTML .....	1587
.....	1589
JVM Summary - HTML.....	1589
.....	1590
JVM System Properties - HTML .....	1590
.....	1591
JVM GC Trends - HTML.....	1591
.....	1592

Tomcat Displays .....	1593
Tomcat Servers .....	1593
All Tomcat Servers .....	1593
Tomcat Server Summary .....	1595
Tomcat Applications .....	1598
Applications Heatmap .....	1599
Applications Summary.....	1600
Tomcat Displays - HTML.....	1603
Tomcat Overview - HTML.....	1603
Tomcat Servers Heatmap - HTML .....	1604
Single Tomcat Server - HTML .....	1605
.....	1606
All Tomcat Apps - HTML .....	1607
.....	1608
Single Tomcat App - HTML.....	1608
<b>Appendix A - Alert Definitions .....</b>	<b>1611</b>
Amazon Web Services .....	1611
Apache Kafka .....	1613
Docker .....	1616
Microsoft SQL Server.....	1618
MongoDB.....	1619
MySQL Database .....	1620
Node.js .....	1621
Oracle Coherence.....	1622
Oracle Database .....	1629
Oracle WebLogic .....	1631
RTView Host Agent.....	1635
RTView Manager and RTView Rules .....	1637
Solace.....	1638
TIBCO ActiveMatrix BusinessWorks.....	1645
TIBCO ActiveSpaces .....	1650
TIBCO ActiveSpaces (2.x) .....	1653
TIBCO Adapters.....	1655
TIBCO BusinessEvents.....	1656
TIBCO Enterprise Message Service .....	1658
TIBCO FTL .....	1663
UX .....	1665
VMware vCenter .....	1666

<b>Appendix B - Limitations</b> .....	1669
iPad Safari Limitations .....	1669
TIBCO ActiveMatrix BusinessWorks.....	1671
Servers.....	1671
Business Works 5.7.1 Engine Status .....	1671
BWSE Components .....	1672
 <b>Appendix C - Third Party Notice Requirements</b> .....	 1673
RTView Enterprise.....	1673
RTView Core .....	1679



# Preface

Welcome to the *RTView Enterprise User's Guide*. Read this preface for an overview of the information provided in this guide and the documentation conventions used throughout, additional reading, and contact information. This preface includes the following sections:

- ["About This Guide"](#)
- ["Additional Resources"](#)
- ["Contacting SL"](#)

---

## About This Guide

The *RTView Enterprise User's Guide* describes how to use RTView Enterprise. For information about how to install and configure RTView Enterprise, see the *RTView Enterprise Configuration Guide*.

### Audience

This guide is written for support teams and managers who are responsible for the performance and health of network and software resources in an organization.

### Document Conventions

This guide uses the following standard set of typographical conventions.

Convention	Meaning
<i>italics</i>	Within text, new terms and emphasized words appear in italic typeface.
<b>boldface</b>	Within text, directory paths, file names, commands and GUI controls appear in bold typeface.
Courier	Code examples appear in Courier font: <pre>amnesiac &gt; enable amnesiac # configure terminal</pre>
< >	Values that you specify appear in angle brackets: <b>interface &lt;ipaddress&gt;</b>

---

## Additional Resources

This section describes resources that supplement the information in this guide. It includes the following information:

- ["Release Notes"](#)
- ["Support Knowledge Base"](#)
- ["SL Documentation"](#)

### Release Notes

The following online file supplements the information in this user guide. It is available on the SL Technical Support site at <http://www.sl.com/support/>.

Examine the online release notes before you begin the installation and configuration process. They contain important information about this release of RTView Enterprise.

### Support Knowledge Base

The SL Knowledge Base is a database of known issues, how-to documents, system requirements, and common error messages. You can browse titles or search for keywords and strings. To access the SL Knowledge Base, log in to the SL Support site located at <http://www.sl.com/support/>.

### SL Documentation

For a complete list and the most current version of SL documentation, visit the SL Support Web site located at [http://www.sl.com/services/support\\_rtviewdocs.shtml](http://www.sl.com/services/support_rtviewdocs.shtml).

For details about configuring RTView Enterprise components, including RTView DataServers, RTView DataCollectors and solution packages, use the following links:

- ["Documentation for IBM and RTView Enterprise"](#)
- ["Documentation for Infrastructure and RTView Enterprise"](#)
- ["Documentation for Kafka and RTView Enterprise"](#)
- ["Documentation for Oracle and RTView Enterprise"](#)
- ["Documentation for Solace and RTView Enterprise"](#)
- ["Documentation for TIBCO and RTView Enterprise"](#)

### Documentation for IBM and RTView Enterprise

- RTView DataServer for IBM Quick Start Guide, Version 5.1.1 (On-Premise): <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/748061094/RTView+DataServer+for+IBM+Quick+Start+Guide+Version+5.1.1+On-Premise>

#### **Solution Packages Documents associated with RTView DataServer for IBM Quick Start Guide:**

- DB2: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/748519712/Configuring+IBM+DB2+Version+5.1.1>
- MQ: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/748453945/Configuring+IBM+WebSphere+MQ+Version+5.1.1>

- WebSphere: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/748519814/Configuring+IBM+WebSphere+Version+5.1.1>
- RTView DataCollector for IBM Quick Start Guide, Version 5.1.1 (On-Premise): <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/747831361/RTView+DataCollector+for+IBM+Quick+Start+Guide+Version+5.1.1+On-Premise>

**Solution Package Documents associated with RTView DataCollector for IBM Quick Start Guide:**

- DB2: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/749043772/Configuring+IBM+DB2+for+DataCollector+Version+5.1.1>
- MQ: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/748847164/Configuring+IBM+WebSphere+MQ+for+DataCollector+Version+5.1.1>
- WebSphere: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/748978345/Configuring+IBM+WebSphere+for+DataCollector+Version+5.1.1>

## Documentation for Infrastructure and RTView Enterprise

- RTView DataServer for Infrastructure (Infra) Quick Start Guide, Version 5.1.1 (On-Premise): <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/747896853/RTView+DataServer+for+Infrastructure+Quick+Start+Guide+Version+5.1.1+On-Premise>

**Solution Package Documents associated with RTView DataServer for Infra Quick Start Guide:**

- AWS: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/748257318/Configuring+Amazon+Web+Services+Version+5.1.1>
- Docker: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/748486673/Configuring+Docker+Version+5.1.1>
- Microsoft SQL: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/748486724/Configuring+Microsoft+SQL+Server+Version+5.1.1>
- MongoDB: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/748519461/Configuring+MongoDB+Version+5.1.1>
- MySQL: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/748519523/Configuring+MySQL+Database+Version+5.1.1>
- Node.js: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/748617737/Configuring+Node.js+Version+5.1.1>
- RedHat JBoss: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/748290188/Configuring+RedHat+JBoss+Version+5.1.1>
- RTView Host Agent: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/748486790/Configuring+RTView+Host+Agent+Version+5.1.1>
- RTView Manager: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/748585004/Configuring+RTView+Manager+Version+5.1.1>
- VMWare: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/748486837/Configuring+VMWare+Version+5.1.1>
- RTView DataCollector for Infrastructure Quick Start Guide, Version 5.1.1 (On-Premise): <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/747929881/RTView+DataCollector+for+Infrastructure+Quick+Start+Guide+Version+5.1.1+On-Premise>

**Solution Packages Documents associated with RTView DataCollector for Infra Quick Start Guide:**

- AWS: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/748978177/Configuring+Amazon+Web+Services+for+DataCollector+Version+5.1.1>
- Docker: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/749043713/Configuring+Docker+for+Data+Collector+Version+5.1.1>
- Microsoft SQL: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/749076481/Configuring+Microsoft+SQL+Server+for+DataCollector+Version+5.1.1>
- MongoDB: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/749076481/Configuring+Microsoft+SQL+Server+for+DataCollector+Version+5.1.1>
- MySQL: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/748978259/Configuring+MySQL+Database+for+DataCollector+Version+5.1.1>
- Node.js: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/748879896/Configuring+Node.js+for+DataCollector+Version+5.1.1>
- RedHat JBoss: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/748879947/Configuring+RedHat+JBoss+for+DataCollector+Version+5.1.1>
- RTView Manager: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/748879947/Configuring+RedHat+JBoss+for+DataCollector+Version+5.1.1>
- VMWare: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/749010964/Configuring+VMWare+for+DataCollector+Version+5.1.1>

**Documentation for Kafka and RTView Enterprise**

- RTView DataServer for Kafka Quick Start Guide, Version 5.1.1 (On-Premise): <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/747766017/RTView+DataServer+for+Kafka+Quick+Start+Guide+Version+5.1.1+On-Premise>

**Solution Package Documents associated with RTView DataServer for Kafka Quick Start Guide:**

- Kafka: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/748552193/Configuring+Apache+Kafka+Version+5.1.1>
- RTView DataCollector for Kafka Quick Start Guide, Version 5.1.1 (On-Premise): <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/748060847/RTView+DataCollector+for+Kafka+Quick+Start+Guide+Version+5.1.1+On-Premise>

**Solution Packages Documents associated with RTView DataCollector for Kafka Quick Start Guide:**

- Kafka: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/748520152/Configuring+Apache+Kafka+for+DataCollector+Version+5.1.1>

**Documentation for Oracle and RTView Enterprise**

- RTView DataServer for Oracle Quick Start Guide, Version 5.1.1 (On-Premise): <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/747798837/RTView+DataServer+for+Oracle+Quick+Start+Guide+Version+5.1.1+On-Premise>

**Solution Package Documents associated with RTView DataServer for Oracle Quick Start Guide:**

- Oracle WebLogic: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/748552829/Configuring+Oracle+WebLogic+Version+5.1.1>

- Oracle Database: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/748585275/Configuring+Oracle+Database+Version+5.1.1>
- RTView DataCollector for Oracle Quick Start Guide, Version 5.1.1 (On-Premise): <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/747962466/RTView+DataCollector+for+Oracle+Quick+Start+Guide+Version+5.1.1+On-Premise>  
**Solution Packages Documents associated with RTView DataCollector for Oracle Quick Start Guide:**
  - Oracle WebLogic: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/748520290/Configuring+Oracle+WebLogic+for+DataCollector+Version+5.1.1>
  - Oracle Database: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/748553402/Configuring+Oracle+Database+for+DataCollector+Version+5.1.1>

## Documentation for Solace and RTView Enterprise

- RTView DataServer for Solace Quick Start Guide, Version 5.1.1 (On-Premise): <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/747929626/RTView+DataServer+for+Solace+Quick+Start+Guide+Version+5.1.1+On-Premise>  
**Solution Packages Documents associated with RTView DataServer for Solace Quick Start Guide:**
  - Solace: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/748323149/Configuring+Solution+Package+for+Solace+Version+5.1.1>
- RTView DataCollector for Solace Quick Start Guide, Version 5.1.1 (On-Premise): <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/747798561/RTView+DataCollector+for+Solace+Quick+Start+Guide+Version+5.1.1+On-Premise>  
**Solution Package Documents associated with RTView DataCollector for Solace Quick Start Guide:**
  - Solace: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/748912641/Configuring+Solution+Package+for+Solace+for+DataCollector+Version+5.1.1>

## Documentation for TIBCO and RTView Enterprise

- RTView DataServer for TIBCO Quick Start Guide, Version 5.1.1 (On-Premise): <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/743047169/RTView+DataServer+for+TIBCO+Quick+Start+Guide+Version+5.1.1+On-Premise>  
**Solution Packages Documents associated with RTView DataServer for TIBCO Quick Start Guide:**
  - BW: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/735117420/Configuring+TIBCO+BusinessWorks+Version+5.1.1>
  - ActiveSpaces 2: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/742686725/Configuring+TIBCO+ActiveSpaces+2+Version+5.1.1>
  - ActiveSpaces: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/742785025/Configuring+TIBCO+ActiveSpaces+Version+5.1.1>
  - ActiveMatrix: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/748322867/Configuring+TIBCO+ActiveMatrix+Version+5.1.1>
  - Adapters: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/748519622/Configuring+TIBCO+Adapters+Version+5.1.1>

- BusinessEvents: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/748290281/Configuring+TIBCO+BusinessEvents+Version+5.1.1>
- EMSMON: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/748552496/Configuring+TIBCO+Enterprise+Message+Service+Version+5.1.1>
- FTL: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/748453894/Configuring+TIBCO+FTL+Version+5.1.1>
- RTView DataCollector for TIBCO Quick Start Guide, Version 5.1.1 (On-Premise): <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/747765793/RTView+DataCollector+for+TIBCO+Quick+Start+Guide+Version+5.1.1+On-Premise>

**Solution Package Documents associated with RTView DataCollector for TIBCO Quick Start Guide:**

- BW: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/748323306/Configuring+TIBCO+BusinessWorks+for+DataCollector+Version+5.1.1>
- ActiveSpaces2: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/748454166/Configuring+TIBCO+ActiveSpaces+2+for+DataCollector+Version+5.1.1>
- ActiveSpaces: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/749830145/Configuring+TIBCO+ActiveSpaces+for+DataCollector+Version+5.1.1>
- ActiveMatrix: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/748454044/Configuring+TIBCO+ActiveMatrix+for+DataCollector+Version+5.1.1>
- Adapters: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/748585408/Configuring+TIBCO+Adapters+for+Data+Collector+Version+5.1.1>
- BE: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/748454257/Configuring+TIBCO+BusinessEvents+for+DataCollector+Version+5.1.1>
- EMSMON: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/748257891/Configuring+TIBCO+Enterprise+Message+Service+for+DataCollector+Version+5.1.1>
- FTL: <https://slcorp.atlassian.net/wiki/spaces/RCS/pages/748847105/Configuring+TIBCO+FTL+for+DataCollector+Version+5.1.1>

---

## Contacting SL

This section describes how to contact departments within SL.

### Internet

You can learn about SL products at <http://www.sl.com>.

## **Technical Support**

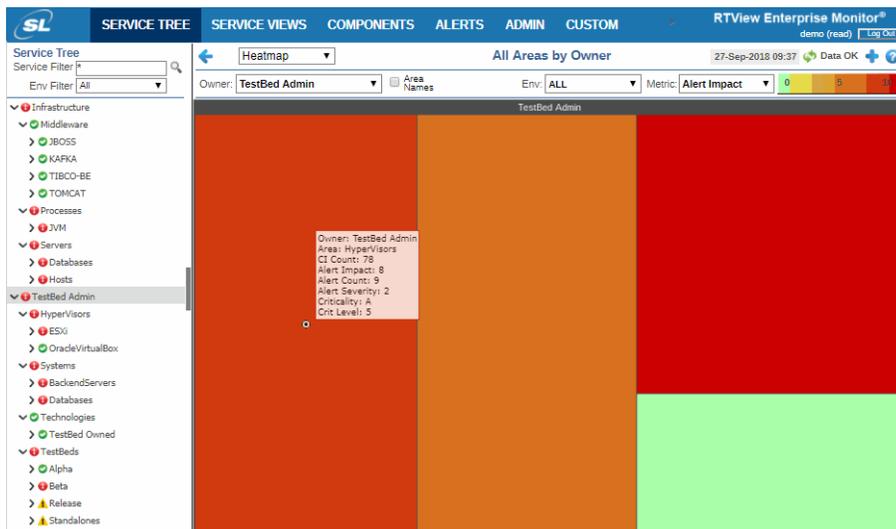
If you have problems installing, using, or replacing SL products, contact SL Support or your channel partner who provides support. To contact SL Support, open a trouble ticket by calling 415 927 8400 in the United States and Canada or +1 415 927 8400 outside the United States.

You can also go to <http://www.sl.com/support/>



# CHAPTER 1 Introduction to RTView Enterprise

Welcome to RTView® Enterprise, a monitoring system that provides single-pane-of-glass visibility of aggregated real-time and historical information about the performance of complex multi-tier applications, including custom-built applications.



For details about setting up RTView Enterprise and RTViewCentral, as well as RTView DataServers and solution packages, see the *RTView Enterprise Configuration Guide*.

This section contains:

- “Displays”
- “System Requirements”

RTView Enterprise has the ability to drill-down to the software-component level to help you determine the root cause of issues affecting application performance. RTView Enterprise (the *Monitor*) enables you to answer questions such as: Are any resources reaching a state of critical health? Do I need to allocate more memory to any resources? Are any having slow response times? Are application deadlocks causing bottlenecks anywhere? Is processing and connection load evenly distributed across resources?

RTView Enterprise enables application support teams to:

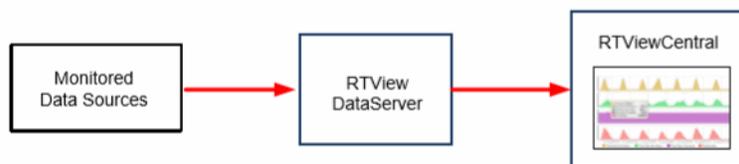
- Provide a single, real-time interface to the end-to-end performance of complex or distributed applications.
- Provide early warning of issues and automate corrective actions tied to alerts, to reduce the number of trouble tickets.
- Leverage historical trends to anticipate possible application degradation and enable preventive care.
- Quickly pinpoint the root cause of issues and initiate repair.
- Reduce costs and minimize lost revenue related to system downtime and degradation.
- Improve performance against SLAs, customer expectations and brand promises.
- Improve business decisions that are tied to application performance.
- Lower the total cost of managing applications.

---

## Displays

RTView® Enterprise provides multiple sets of displays for monitoring your system. Some displays come with and reside on “RTViewCentral”. Additional displays can be added on via “RTView DataServers”. This section describes RTViewCentral and RTView DataServers, their roles and the types of displays they provide.

The following figure illustrates “RTViewCentral”, a single “RTView DataServers” and the basic data flow from the monitored data sources.



RTView DataServers collect and store metric data from your data sources. “RTViewCentral” provides the graphic visualization of the metric data collected by RTView DataServers. Performance data collected by RTView DataServers are correlated with the displays that come with RTView Enterprise.

### RTViewCentral

RTViewCentral is where the metric data collected by the RTView DataServer is analyzed, correlated and transformed; historical data is aggregated; alert rules and actions are defined; and where the “master” mapping of everything monitored in your system resides.

RTViewCentral is comprised of a Display Server, the Central Server, the Central Alert Historian and a database. RTViewCentral also has “RTView Manager” displays and RTView RTRules.

RTViewCentral displays come with RTView Enterprise and reside on RTViewCentral.

Find these displays organized under the following Views:

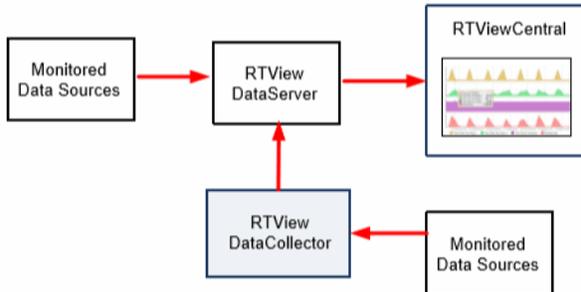
- **"All Management Areas"**: These displays show the health of your entire system using aggregated data from all Areas. Use these displays to quickly identify critical conditions across all Areas in your system, then drill-down to investigate in lower-level displays.
- **"Multi Area Service Views"**: These displays show the health of Services for one or more Groups. Use these displays to identify critical conditions across all Areas or a single Area. Drill-down to investigate in lower-level displays.
- **"Single Area Service Views"**: These displays show the health of Services for one or more Groups. Use these displays to identify critical conditions across a single Area. Drill-down to investigate in lower-level displays.
- **"Service Summary Views"**: These displays show the health of CI Types. Use these displays for a closer view of a critical condition, including alert details.
- **"Key Metrics Views"**: These displays show how close a metric is approaching its threshold (rather than your ACTIVE alerts and their impact on the overall application or service), enabling you to anticipate performance problems BEFORE the alert threshold is crossed and analyze the circumstances that led up to error conditions.
- **"Component Views"**: These displays show the lowest level view of CMDB contents--the component level. In these displays, alert states for components are shown by Service and Area in tabular and heatmap formats, while highlighting the most critical alert state for each component.
- **"Metric Explorer"**: The Metric Explorer (MX) is a tool for creating and viewing custom dashboards, referred to as MX Views. An MX View contains a trend graph with up to five traces which you can configure to show numeric metrics from any solution package.
- **"Alert Views"**: These displays show detailed information about all alerts that have occurred in your RTView Enterprise system (all Owners and all Areas).
- **"Administration"**: These displays enable you to set alert thresholds, and observe how alerts are managed.
- **"CMDB Administration"**: This display allows you to modify your Service Data Model.
- **"Architecture"**: These displays provide a view of RTView Enterprise component connectivity, mapping between component types, and component level connection and performance information.
- **"Property Views"**: These displays show how your Monitor properties are configured and the values for all connected RTView processes.
- **"Diagram Views"**: These displays are dedicated for the Diagram Generator, a feature that auto-creates diagram displays which mirror your system components and hierarchy.

## RTView DataServers

RTView DataServer displays reside on the RTView DataServer. You can add displays to your RTView Enterprise system by installing one or more RTView DataServers. RTView DataServers contain a series of displays based on the type of technology being monitored. **"Solution Packages"** are bundled into RTView DataServers. For example, the RTView DataServer for TIBCO, which includes the Solution Packages for TIBCO EMS, BusinessWorks, and many other TIBCO applications.

RTView DataServers are also available for IBM, Infrastructure, Kafka, Oracle, Solace and RTView Manager. Performance data collected for these technologies are correlated with the displays that come with RTView Enterprise.

Each RTView DataServer has a corresponding *RTView DataCollector*. The RTView DataCollector is useful for distributed deployments as they can be deployed to collect data from otherwise unreachable data sources and send the data to its RTView DataServer pair. The following figure illustrates RTView DataCollector deployment.



The following RTView DataServers are available:

- **"RTView DataServer for IBM"**: Used to monitor the health and performance across all components for IBM® MQ, IBM DB2 databases and IBM WebSphere servers and applications.
- **"RTView DataServer for Infrastructure"**: Used to monitor the health and performance across all components for Amazon Web Services, Docker, JBoss, MongoDB, MySQL, MS SQL Server, Node.js, VMware, UX (User Experience), and RTView Host Agents.
- **"RTView DataServer for Kafka"**: Used to monitor the health and performance across all components for Apache Kafka.
- **"RTView DataServer for Oracle"**: Used to monitor the health and performance across all components for Oracle Coherence, Oracle Database, Connector for Oracle Enterprise Manager and Oracle WebLogic.
- **"RTView DataServer for Solace"**: Used to monitor the health and performance across all components for Solace routers, bridges, endpoints, clients and Syslog events.
- **"RTView DataServer for TIBCO"**: Used to monitor the health and performance of TIBCO ActiveMatrix, TIBCO ActiveSpaces, TIBCO Adapters, TIBCO BusinessEvents, TIBCO BusinessWorks, TIBCO Enterprise Message Service, TIBCO FTL, and TIBCO Hawk.

## Solution Packages

Solution packages gather metrics from infrastructure, middleware, instrumented applications, JVMs, log files, and third party monitoring products. RTView Enterprise also provides a means for creating custom solution packages to gather most any piece of performance information with a wide array of built-in data adapters. These custom solution packages can be configured without programming. SL Support has many templates for custom solution packages that can be delivered to users or customized as a service.

A solution package provides these main pieces of functionality to RTView Enterprise:

- **Data Access:** The solution package gathers the performance metrics relevant to the technology being monitored. The data may be gathered by either synchronous or asynchronous direct connections to a technology, or by receiving information from RTView agents deployed on the hosts of the monitored technology.
- **Data Caching:** Performance metrics are stored in in-memory data caches to supply quick access to the most current performance metrics.
- **Data History:** Long-term performance metrics can be stored in a JDBC-enabled relational database. The solution package allows for the configuration of the rules for data compaction and management of long-term data persistence.
- **Alert Event Access:** If the solution package is connecting to another monitoring system, it can gather alert events from that system, bring those events into RTViewCentral and allow alert management to be performed in RTViewCentral. Optionally, the solution package can be configured to synchronize alert states between the two systems.
- **Alert Rules Engine:** The solution packages are configured with alert rule definitions which are processed real-time in the RTView DataServers. Dynamic updates to these alert rule definitions, such as changing alert rule thresholds or policies, can be managed through the RTViewCentral **Alert Administration** interface. When alerts are activated by these alert rule definitions, they are sent to RTViewCentral to be aggregated with other solution package alerts.
- **Data Viewing:** Each solution package comes with designated displays which can be accessed by RTView Enterprise to show the performance metrics in summary and drill-down views.
- **Data Server:** This Java process is run to begin accessing the data, storing data to internal memory caches, running the alert rules and optionally providing data to the Historian process.
- **Data Historian:** The process manages the storage of information into a relational database and runs the rules relevant to managing this persisted data.

---

## System Requirements

For browser support, hardware requirements, JVM support and other system requirement information, please refer to the **README\_sysreq.txt** file from your product installation. A copy of this file is also available on the product download page.



## CHAPTER 2 Using the Monitor

Welcome to RTView Enterprise. This section describes how to access the Monitor and the many displays that come with RTViewCentral. There are two versions of RTView Enterprise, the Display Server version and the HTML version. The Display Server version is described unless otherwise noted.

This section contains:

- [“Login to RTView Enterprise”](#)
- [“Overview”](#): Describes the user interface and navigation, behavior of graphic objects, color codes and icons.
- [“RTViewCentral Displays”](#): Describes the displays that come with RTView Enterprise.
- [“Other - HTML Displays”](#): Describes [“Common/Alerts - HTML”](#) displays and [“Common/System - HTML”](#) displays.

For an overview and details about configuring RTView Enterprise, including RTViewCentral, RTView DataServers, RTView DataCollectors and solution packages, see the *RTView Enterprise Configuration Guide*.

---

## Login to RTView Enterprise

This section contains:

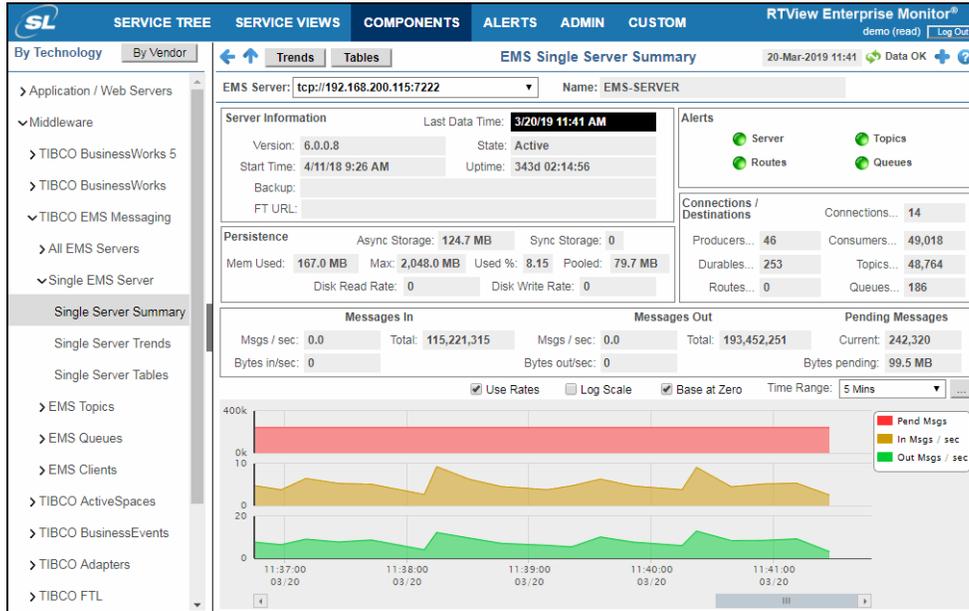
- [“Display Server Login”](#)
- [“HTML Login”](#)

### Display Server Login

#### To access RTView Enterprise:

- Set the JAVA\_HOME environment variable to the location of your Java installation.
- Execute the **start\_servers** script, located in the **RTViewCentral/bin** directory, to start RTView Enterprise.
- Browse to **http://localhost:10070/rtview-central-classic** and login (use admin/admin)

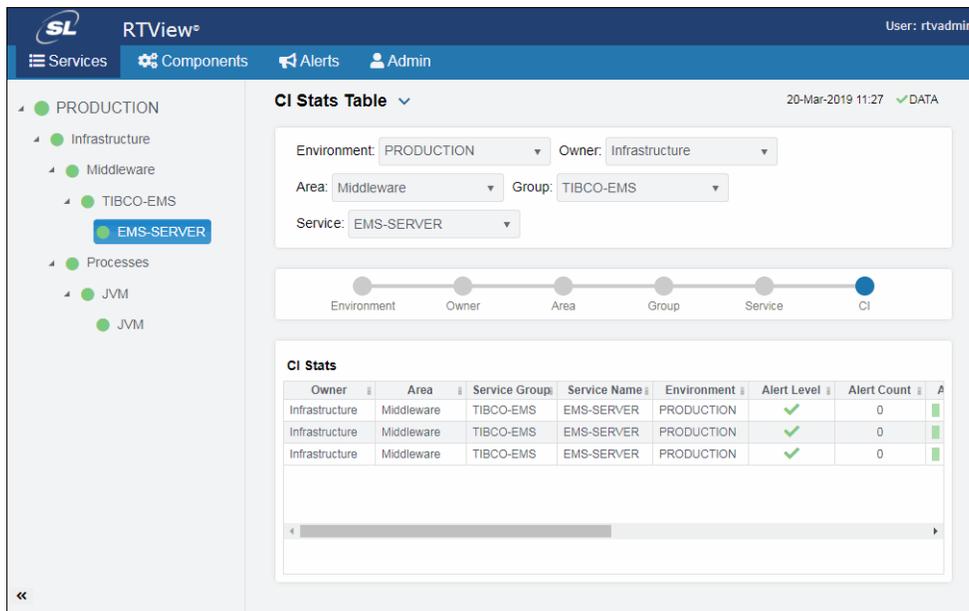
By default, data is collected every 15 seconds and displays are refreshed 15 seconds afterward.



## HTML Login

To open the HTML version of RTView Enterprise browse to: **http://localhost:11070/rtview-central** (login as rtvadmin/rtvadmin or rtvuser/rtvuser).

By default, data is collected every 15 seconds and displays are refreshed 15 seconds afterward.



For details about displays that you can add-on via RTView DataServers, see the following:

- ["RTView DataServer for IBM"](#)
- ["RTView DataServer for Infrastructure"](#)
- ["RTView DataServer for Kafka"](#)
- ["RTView DataServer for Oracle"](#)
- ["RTView DataServer for Solace"](#)
- ["RTView DataServer for TIBCO"](#)

---

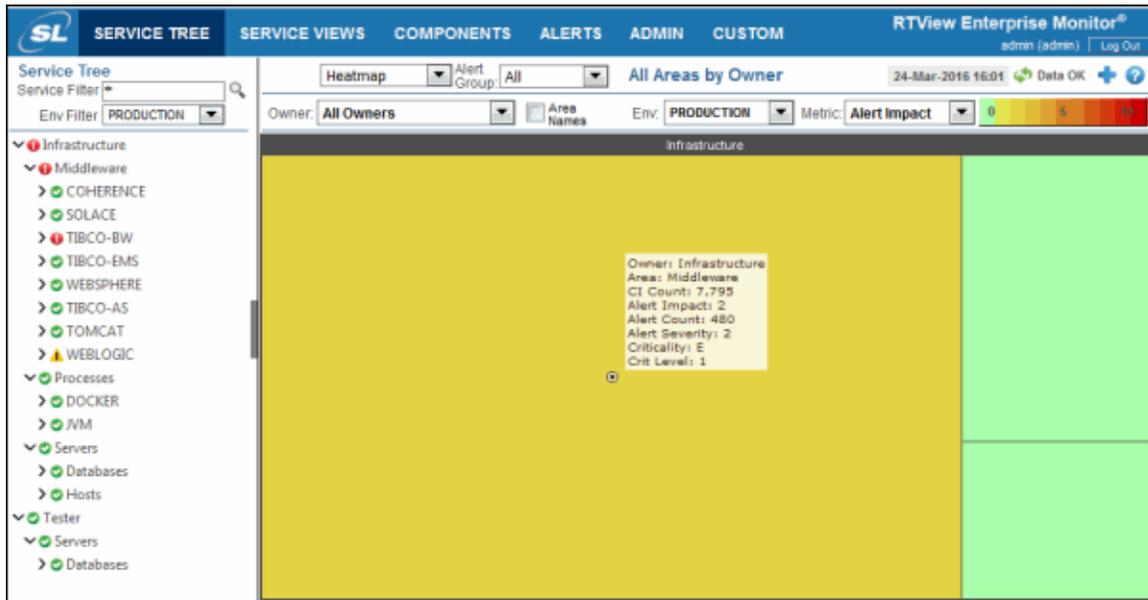
## Overview

RTView Enterprise uses visual cues (such as color coding, graphic charts and sizing of shapes) to communicate the current state of all elements in your system. This section describes how displays are structured and organized, how to read heatmaps, tables and trend graphs, as well as GUI functionality and navigation in the Display Server version.

This section includes:

- ["Navigation Tabs"](#)
- ["Fundamental Structure of Displays"](#)
- ["Heatmaps"](#)
- ["History Heatmaps"](#)
- ["Tables"](#)
- ["Trend Graphs"](#)
- ["Popup Menu"](#)
- ["Title Bar"](#)

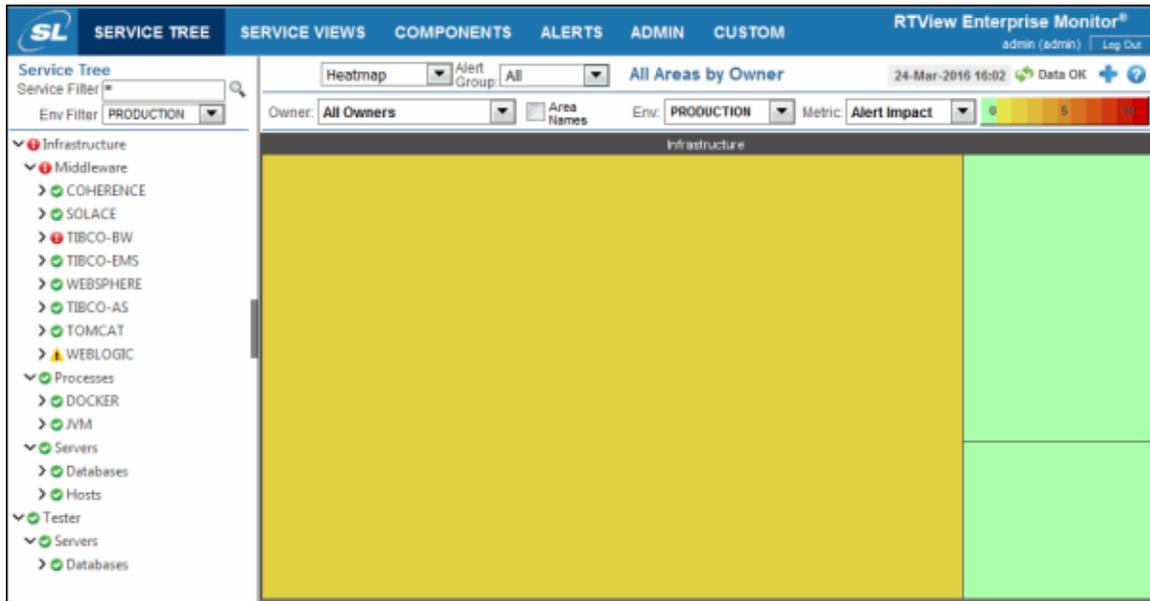
## Navigation Tabs



There are six different tabs that run along the upper portion of the window:

- **"SERVICE TREE Tab"**: provides a tree structure view of your defined CMDB with 4 levels of hierarchy: **Owner>Area>Group>Service**. The tree is configurable and should be set up to represent how a support person likes to conceptually think of the vast number of items that are being monitored.
- **"SERVICE VIEWS Tab"**: provides an alternate way of accessing the primary RTView Enterprise displays also found in the **SERVICE TREE** tab. This tab might be used by power users who are very familiar with their monitoring environment and choose not to visualize the entire service tree.
- **"COMPONENTS Tab"**: provides access to the **JVM Process Views**, the **Tomcat Servers Views**, the **RTView Servers Views**, and any Views included with the solution packages that you have installed. This tab organizes the monitoring information by technology or vendor and allows you to view the health state of your technology footprint without logical or service groupings. Specialists that understand in some depth how the technologies are deployed, configured, load-balanced, and scale based on load will gain benefit on the organization of performance metrics by the important functional sub-components of the technology.
- **"ALERTS Tab"**: provides a view of the current active alerts in the system and allows you to manage those alerts by owning them, acknowledging them, and/or suppressing them.
- **"ADMIN Tab"**: can be accessed by administrators of RTView Enterprise, who can use this tab during installation to set up proper alert settings, to describe logical and service groupings that drive the construction of the Service Tree, and to "monitor the monitor" view of the current health state of RTView Enterprise and how it is currently deployed and configured.
- **"CUSTOM Tab"**: provides a location where you can add your own tab and views.

## SERVICE TREE Tab



The **SERVICE TREE** tab provides a tree structure view of your defined CMDB with 4 levels of hierarchy: **Owner>Area>Group>Service** (see [“Fundamental Structure of Displays”](#) for more information). This tab is the primary source for understanding the health of your services and for drilling down to analyze issues. The Service Tree, which is configurable, shows user-defined logical groupings of the infrastructure and middleware used to support applications and should be set up to represent how support personnel like to conceptually think of the vast number of items that are being monitored. These groupings could, for example, contain a collection of monitored component ID (CIs) used to support a specific application or a service, or they could contain CIs relevant to varying technologies located at specific data centers. The Service Tree aggregates the current alert state of any item in a group to indicate which groups have items that need to be investigated, and you can use a variety of visual clues to prioritize and analyze the issues. You can also determine priority using the Alert Impact view in the heatmaps to identify which alert conditions will be the most impactful to your business, and you can then analyze the situation using a variety of tools including:

- **Key Metrics:** allows you to view the cross-correlation of CIs relevant to a grouping or service and how their performance may affect each other and the services they support. For details, see [“Key Metrics Views”](#).
- **Drill Down CI Summary Views:** provides a way to analyze how a particular CI has been performing over time.
- **Metric Explorer:** allows you to choose specific metrics to chart when analyzing several critical performance metrics over time. For details, see [“Metric Explorer”](#).

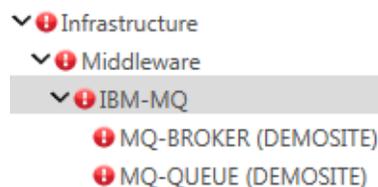
This tab allows you to filter the navigation tree content by service and environment (see figure below). The environment you select also sets the **Environment** filter on the main panel. Note that changing the **Environment** filter on the main panel does not set the **Environment** filter in the navigation panel.



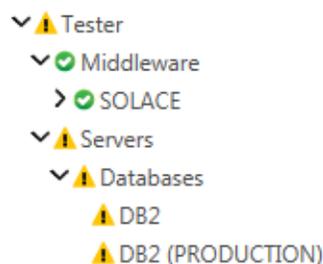
## Alerts

Each level within the Service Tree has a red, yellow, or green icon next to it, which indicate the highest alert level for that particular Owner, Area, Group, or Service. These icons allow you to instantly recognize problem areas within your system and allow you to drill down to quickly find the source of the issue. A red icon  indicates that one or more alerts exceeded their ALARM LEVEL threshold, a yellow icon  indicates that one or more alerts exceeded their WARNING LEVEL threshold, and a green icon  indicates that no alerts exceeded their WARNING or ALARM LEVEL threshold.

The Owner, Area, and Group automatically display the icon representing the highest level alert for their associated services with red (ALARM LEVEL threshold exceeded) being the most serious, yellow (WARNING LEVEL threshold exceeded) being intermediate, and green meaning everything is functioning normally. For example, if any of the services within a particular **Owner>Area>Group** have one or more alerts that exceeded their ALARM LEVEL threshold and, hence, have a red icon next to it in the tree, then the associated Owner, Area, and Group levels will also have the same red icon. In the example below, you can see that the MQ Broker service has one or more alerts that exceeded their ALARM LEVEL threshold and has a red indicator. As a result, the Owner, Area, and Group also have the red indicator



If the highest alert level for the services within a particular **Owner>Area>Group** is a service that has one or more alerts that exceeded their WARNING LEVEL threshold and, hence, has a yellow icon next to it in the tree, then the associated Owner, Area, and Group levels will also have the same yellow icon. In the example below, you can see that the DB2 database has one or more alerts that exceeded its WARNING LEVEL threshold and has a yellow indicator. Since none of the other services in this particular tree have alerts that exceeded their ALARM LEVEL threshold, then the associated Owner, Area, and Group also have the yellow indicator since the WARNING LEVEL threshold is the highest alert level threshold exceeded.



## Available Displays

The following displays are available in the following levels in this tab:

### Owner Level (top level)

To access the following displays, select an Owner Level option (**Infrastructure**, for example) to display an Owner level display. Select one of the following options from the drop-down in the upper left-hand corner of the display to view the associated display:

Drop-down Option	Display	Description
Heatmap	"Area Heatmap"	Heatmap of the most critical alerts for all Areas of your system, with the option to filter by Owner, Environment and alert Metric.
Area	"Area Table"	Table of data shown in the "Area Heatmap" with the option to filter by Owner and Environment.

**Note:** When selecting an Owner Level option, the display that opens by default will be the one that was last selected. For example, if Heatmap was the display that was previously selected, Heatmap will display by default again.

### Area Level (second level down)

To access the following displays, select an Area Level option (**Middleware** in the example above) to display an Area Level display. Select one of the following options from the drop-down in the upper left-hand corner of the display to view the associated display:

Drop-down Option	Display	Description
By Group	"Group/Service Heatmap"	Heatmap of alert states for Services by Area, with the option to filter by Area, Group, Environment and alert Metric, and the option to show Group and Service Names.
By Region	"Group/Region Heatmap"	Heatmap as described for the <b>Group / Service Heatmap</b> (above), with the option to filter by Region and no option to show Service Names.
Table	"Group / Service Table"	Table of data shown in the "Group/Service Heatmap".
By CI Type	"Services CI Type Summary"	Table that shows the health state of Services per CI Type.
History	"Services History Heatmap"	Heatmap of alert states, over time, for Services in a selected Area, with the option to filter by Group, Environment and alert Metric.

**Note:** When selecting an Area Level option, the display that opens by default will be the one that was last selected. For example, if Group/Service Heatmap was the display that was previously selected, Group/Service Heatmap will display by default again.

### Group Level (third level down)

To access the following displays, select a Group Level option (IBM-MQ in the example above) to display a Group Level display. Select one of the following options from the drop-down in the upper left-hand corner of the display to view the associated display:

Drop-down Option	Display	Description
<b>By Group</b>	"Single Area: Group/Service Heatmap"	Heatmap of alert states for Services by Area, with the option to filter by Area, Group, Environment and alert Metric, and the option to show Group and Service Names.
<b>By Region</b>	"Single Area: Region/Service Heatmap"	Heatmap as described for the <b>Group / Service Heatmap</b> (above), with the option to filter by Region and no option to show Service Names.
<b>Table</b>	"Single Area: Region/Service Heatmap"	Table of the data shown in the "Single Area: Group/Service Heatmap".
<b>By CI Type</b>	"Single Area: Services CI Type Summary"	Table that shows the health state of Services per CI Type.
<b>History</b>	"Single Area: Services History Heatmap"	Heatmap of alert states, over time, for Services in a selected Area, with the option to filter by Group, Environment and alert Metric.

---

**Note:** When selecting a Group Level option, the display that opens by default will be the one that was last selected. For example, if Group/Service Heatmap was the display that was previously selected, Group/Service Heatmap will display by default again.

---

### Service Level (fourth level down)

To access the following displays, select a Service Level option (MQ Broker (DEMOSITE) in the example above) to display a Service Level display. Select one of the following options from the drop-down in the upper left-hand corner of the display to view the associated display:

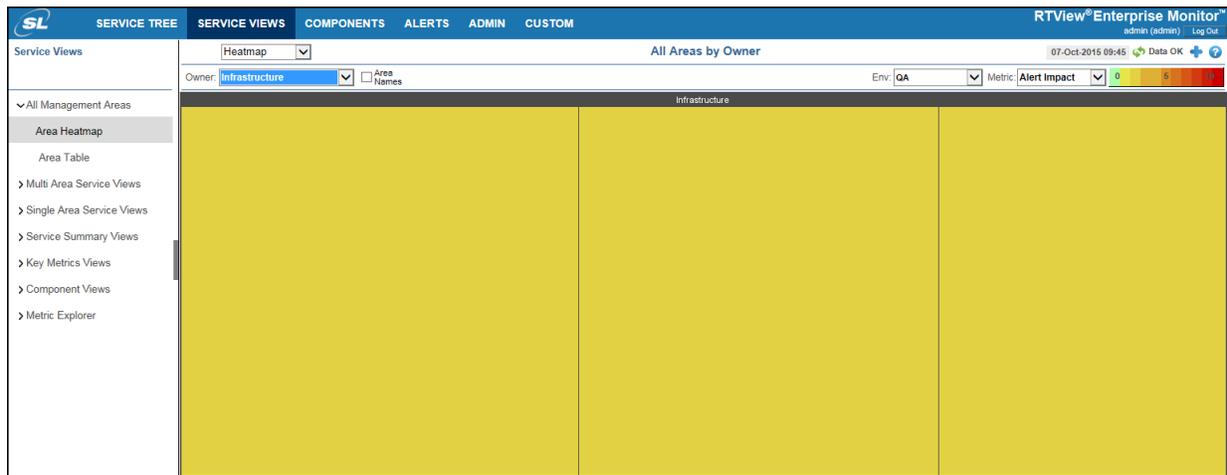
Drop-down Option	Display	Description
<b>By CI Type</b>	"Service By CI Type"	Table of alert states for a Service organized CI Type, with general alert information.
<b>Summary</b>	"Service Summary"	Table of CIs by Service, with detailed alert information.
<b>Health</b>	"Service Health Heatmap"	Heatmap of CIs by Service, with the option to filter by Owner, Area, Group, Environment and alert Metric, and show CI Names.
<b>KM Heatmap</b>	"Service KM Heatmap"	Heatmap of Key Metrics current data for one or more Services in your CMDB hierarchy.
<b>KM Table</b>	"Service KM Table"	Table of Key Metrics current data for one or more Services.
<b>KM History</b>	"Service KM History"	History heatmap of Key Metrics historical data for one or more Services.
<b>KM History (Alt)</b>	"Service KM History (Alt)"	History heatmap of Key Metrics historical data for one or more Services.

**Note:** When selecting a Service Level option, the display that opens by default will be the one that was last selected. For example, if By CI Type was the display that was previously selected, By CI Type will display by default again.

Select the following button, which is available when you select either **By CI Type (Service By CI Type display)** or **Summary (Service Summary display)** from the drop-down list, to open the associated display:

Button	Display	Description
	"Metric Explorer"	The Metric Explorer (MX) is a tool for creating and viewing custom dashboards, referred to as MX Views.

## SERVICE VIEWS Tab



The **SERVICE VIEWS** tab is a simplified version of the **SERVICE TREE** tab that uses drop-down navigation to access displays without the complexity of the service tree. This tab contains the following Views:

- **"All Management Areas"**: Displays in this View show the health of your entire system using aggregated data from all Areas. Use these displays to quickly identify critical conditions across all Areas in your system, then drill-down to investigate in lower-level displays.
- **"Multi Area Service Views"**: Displays in this View show the health of Services for one or more Groups. Use these displays to identify critical conditions across all Areas or a single Area. Drill-down to investigate in lower-level displays.
- **"Single Area Service Views"**: Displays in this View show the health of Services for one or more Groups. Use these displays to identify critical conditions across a single Area. Drill-down to investigate in lower-level displays.
- **"Service Summary Views"**: Displays in this View show the health of CI Types. Use these displays for a closer view of a critical condition, including alert details.
- **"Key Metrics Views"**: The Key Metrics (KM) feature shows how close a metric is approaching its threshold (rather than your ACTIVE alerts and their impact on the overall application or service), enabling you to anticipate performance problems BEFORE the alert threshold is crossed and analyze the circumstances that led up to error conditions.
- **"Metric Explorer"**: The Metric Explorer (MX) is a tool that allows end-users to quickly create custom dashboards for metrics they specifically want to analyze.

## COMPONENTS Tab

The **COMPONENTS** tab organizes the monitoring information by technology or vendor and allows you to view the health state of your technology footprint without logical or service groupings. This tab also contains deep summaries and drill-downs to the subcomponents that comprise a particular technology. By default, this tab provides access to the **JVM Process Views**, the **Tomcat Servers Views**, the **RTView Servers Views**, and any Views included with the solution packages that you have installed.

**NOTE:** The COMPONENTS tab navigation tree (in left panel) has been enhanced to filter the available displays based on the solution packages that are hosted by the RTView DataServer connections enabled by your administrator. If you run a project from a previous release where the navigation tree was customized, this filter is applied in addition to any customization.

The following views are available via this tab:

- **"JVM Processes View/Displays"**: Displays in this View show performance data for monitored Java Virtual Machine (JVM) Processes. Use these displays to monitor performance of your JVMs.
- **"Tomcat Monitor Views/Displays"**: Displays in this View show performance data for monitored Tomcat applications. Use these displays to monitor Tomcat connections and performance of your Web applications and modules.

There are two different ways to view the available displays: **By Technology** and **By Vendor**.

### By Technology Button

The **By Technology** button lists the available displays by the type of technology (Application/ Web Servers, Middleware, Databases, Processes, Hosts/VMs, Connectors, Other).

RTView® Enterprise Monitor™											
admin (admin) Log Out											
All JVMs - Table View											
23-Sep-2015 17:02 Data OK											
JVM Count: 51 <input checked="" type="checkbox"/> Show Inactive											
All JMX Connections											
Connection	Expired	Connected	Alert	Host	Port	CPU %	Max Heap	Mem Used %	Display Name	URL	RtvAppTy...
ALERT_SERVER	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	localhost	10023	18.6	492,896,256	55.7			3 local
ALERTHISTORIAN	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	localhost	10025	0.6	477,233,152	4.1			11 local
AMXMON-HISTORIAN	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	localhost	3367		0				0 local
AMXMON-SLHOST-WIN3	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	192.168.200.133	6368	2.0	954,466,304	37.8			3 local
AMXMON-SLHOST-WIN4	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	192.168.200.134	6368	2.0	954,466,304	31.7			3 local
BW6MON-SLHOST-WIN3	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	192.168.200.133	3368	0.9	954,466,304	20.2			3 local
BW6MON-SLHOST-WIN4	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	192.168.200.134	3368	1.0	954,466,304	20.2			3 local
BWMON-HISTORIAN	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	localhost	3367		0				0 local
BWMONITOR-WIN-8	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	192.168.200.138	3368		0				0 local
CONFIG_SERVER	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	localhost	10013	2.4	477,233,152	34.9			3 local
DISPLAYSERVER	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	localhost	10024	4.0	477,233,152	62.9			5 local
DISPLAYSERVER_DARK...	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	localhost	10124	2.5	477,233,152	29.9			5 local
EMSMON-HISTORIAN	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	localhost	3167		0				0 local
EMSMONITOR-WIN-8	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	192.168.200.138	3168	1.3	954,466,304	28.6			3 local
EMSMON-SLHOST-WIN3	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	192.168.200.133	3168	1.9	954,466,304	17.1			3 local
EMSMON-SLHOST-WIN4	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	192.168.200.134	3168	1.6	954,466,304	20.4			3 local
local	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	localhost		1.8	954,466,304	12.8		local	3 local
MISCMON-HISTORIAN	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	localhost	3967		0				0 local
MISCMON-SLHOST-WIN3	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	192.168.200.133	3968	13.0	1,071,316,992	95.4			3 local
MISCMON-SLHOST-WIN4	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	192.168.200.134	3968	5.3	985,661,440	64.4			3 local
MQMON-64-OL7-3	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	192.168.200.73	3468	4.2	1,037,959,168	9.4			3 local
MQMON-HISTORIAN	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	localhost	3467		0				0 local
MQMON-SLHOST-WIN3	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	192.168.200.133	3468	3.7	954,466,304	35.6			3 local
OCMON-64-OL7-1	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	192.168.200.71	9911		0				0 local
OCMON-64-OL7-4	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	192.168.200.74	9911	0.4	954,728,448	1.6			3 local
OCMONITOR-WIN-8	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	192.168.200.138	9911		0				0 local
OCMON-SLHOST-WIN3	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	192.168.200.133	9911	3.8	954,466,304	27.6			3 local
OCMON-SLHOST-WIN7	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	192.168.200.137	9911		0				0 local
RTVMGR-HISTORIAN	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	localhost	3067		0				0 local
RTVMGR-SLHOST-WIN3	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	192.168.200.133	3068	1.7	954,466,304	10.9			3 local
RTVMGR-SLHOST-WIN4	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	192.168.200.134	3068	1.8	954,466,304	12.8			3 local
RTVRULES	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	192.168.200.134	3868	0.5	715,849,728	10.2			3 local
RTVRULES-SLHOST-WIN3	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	192.168.200.133	3868	0.5	715,849,728	18.1			3 local
SOLMON-64-OL7-6	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	192.168.200.76	4168	1.5	954,728,448	47.7			3 local

### By Vendor Button

The **By Vendor** button lists the displays by vendor name (for example: TIBCO, Oracle, and IBM).

The screenshot shows the RTView Enterprise Monitor interface. The top navigation bar includes 'SERVICE TREE', 'SERVICE VIEWS', 'COMPONENTS', 'ALERTS', 'ADMIN', and 'CUSTOM'. The 'Alerts Table' is active, displaying a list of alerts. The interface includes various filters and controls:

- Alert Group:** Current
- Alert Filter:** All
- Search Text:** (empty)
- CMDB Filter:** Owner = Infrastructure | Area = Middleware | Group = \* | Service = \* | Env = \*
- Alert Settings:** Conn OK
- Summary:** Total 243 / 243, Critical 243 / 243, Warning 0 / 0, Suppressed 0
- Alerts Table Columns:** First Occ, Last Occ, Count, Sup, Owner, Alert Name, Primary Service, CI
- Alerts Table Data:** Multiple rows of alerts, all with 'High Alert' status and 'Bw-PROCESS' as the primary service.
- Columns:** Id, Closed, Closed Reason, Alert Index
- Buttons:** Go To CI, Options, Details

## ALERTS Tab

The screenshot shows the RTView Enterprise Monitor interface. The top navigation bar includes 'SERVICE TREE', 'SERVICE VIEWS', 'COMPONENTS', 'ALERTS', 'ADMIN', and 'CUSTOM'. The 'ALERTS' tab is active. On the left, a service tree is visible with 'Alerts' selected. The main area displays a table of active alerts. The table has columns: Total (197 / 197), Critical (197 / 197), Warning (0 / 0), and Suppressed (0). The table lists multiple instances of 'BwActivityExecutionTimeHi' alerts. The top right shows 'RTView Enterprise Monitor' and '11-Apr-2016 16:00'.

The **ALERTS** tab provides a view of the current active alerts in the system and allows you to manage those alerts by owning them, acknowledging them, and/or suppressing them. You can navigate and filter the alert list by using the service tree to focus on alerts by logical or service groupings. This tab is customizable and can be interfaced with an existing trouble ticket system so that alerts that require an action can be tracked and managed by those systems.

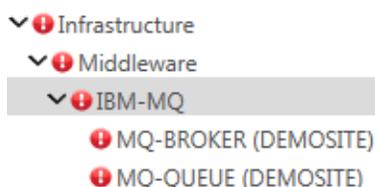
This tab allows you to filter the navigation tree content by service and environment (see figure below). The environment you select also sets the **Environment** filter on the main panel. Note that changing the **Environment** filter on the main panel does not set the **Environment** filter in the navigation panel.



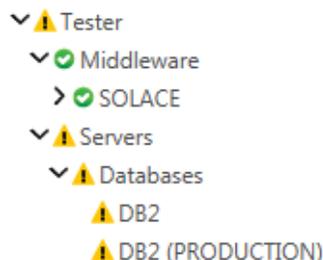
### Alert Icons

Each level within the Alerts tab service tree has a red, yellow, or green icon next to it, which indicate the highest alert level for that particular Owner, Area, Group, or Service. These icons allow you to instantly recognize problem areas within your system and allow you to drill down to quickly find the source of the issue. A red icon  indicates that one or more alerts exceeded their ALARM LEVEL threshold, a yellow icon  indicates that one or more alerts exceeded their WARNING LEVEL threshold, and a green icon  indicates that no alerts exceeded their WARNING or ALARM LEVEL threshold.

The Owner, Area, and Group automatically display the icon representing the highest level alert for their associated services with red (ALARM LEVEL threshold exceeded) being the most serious, yellow (WARNING LEVEL threshold exceeded) being intermediate, and green meaning everything is functioning normally. For example, if any of the services within a particular **Owner>Area>Group** have one or more alerts that exceeded their ALARM LEVEL threshold and, hence, have a red icon next to it in the tree, then the associated Owner, Area, and Group levels will also have the same red icon. In the example below, you can see that the MQ Broker service has one or more alerts that exceeded their ALARM LEVEL threshold and has a red indicator. As a result, the Owner, Area, and Group also have the red indicator



If the highest alert level for the services within a particular **Owner>Area>Group** is a service that has one or more alerts that exceeded their WARNING LEVEL threshold and, hence, has a yellow icon next to it in the tree, then the associated Owner, Area, and Group levels will also have the same yellow icon. In the example below, you can see that the DB2 database has one or more alerts that exceeded its WARNING LEVEL threshold and has a yellow indicator. Since none of the other services in this particular tree have alerts that exceeded their ALARM LEVEL threshold, then the associated Owner, Area, and Group also have the yellow indicator since the WARNING LEVEL threshold is the highest alert level threshold exceeded.



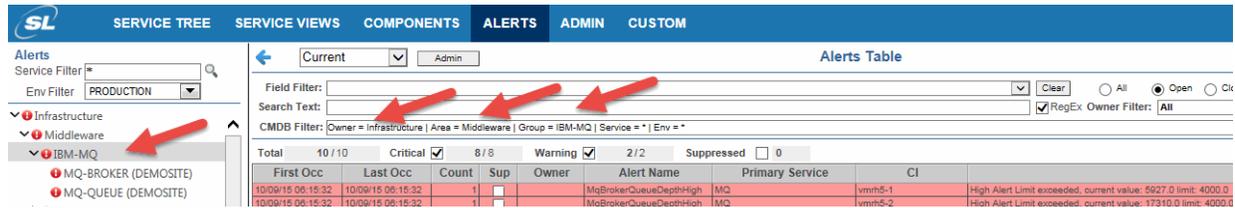
## Available Displays

To access the following displays, select one of the following options from the drop-down in the upper left-hand corner of the display to view the associated display:

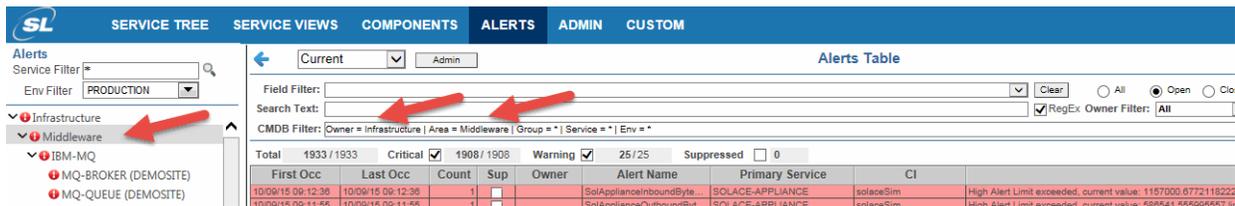
Drop-down Option	Display	Description
<b>Current</b>	"RTView Alerts Table"	This display allows you to track and manage all alerts that have occurred in the system, as well as to add comments, acknowledge, or assign Owners to alerts.
<b>History</b>	""	This display allows you to track the history of any alert that has occurred in your RTView Enterprise system.

**Note:** When selecting an option at any level, the display that opens by default will be the one that was last selected. For example, if History was the display that was previously selected, History will display by default again.

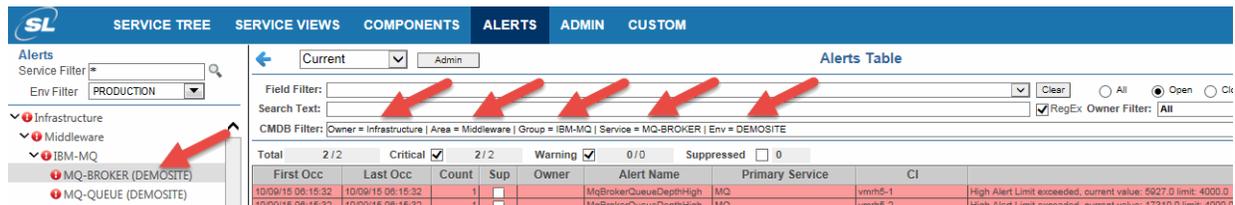
When you select an option at any of the **Owner>Area>Group>Services** levels in the **RTView Alerts Table** display, the display automatically filters the list of alerts based on the level you selected. For example, if you were to select the IBM-MQ option at the Group level, then the filter will be set to Owner=Infrastructure, Area=Middleware, Group=IBM-MQ (the option at the level you selected), and Service and Environment will be set to \* (or all services and environments for that particular Group).



If you were to select the Middleware option at the Area level, then the filter will be set to Owner=Infrastructure, Area=Middleware (the option at the level you selected), and Group and Service and Environment will be set to \* (or all groups, services, and environments for that particular Area).



If you were to select the MQ-BROKER option at the Service level, then the filter will be set to Owner=Infrastructure, Area=Middleware, Group=IBM-MQ, Service=MQ-BROKER, and Environment=DEMOSITE (the option at the level you selected).



### Available Display via a Button

If you select the **Current** option from the drop-down list, the following button is available on the **RTView Alerts Table** display. Select the following button to open the associated display:

Button	Display	Description
	"Alert Administration"	This display allows you to set global or override alert thresholds.

## ADMIN Tab

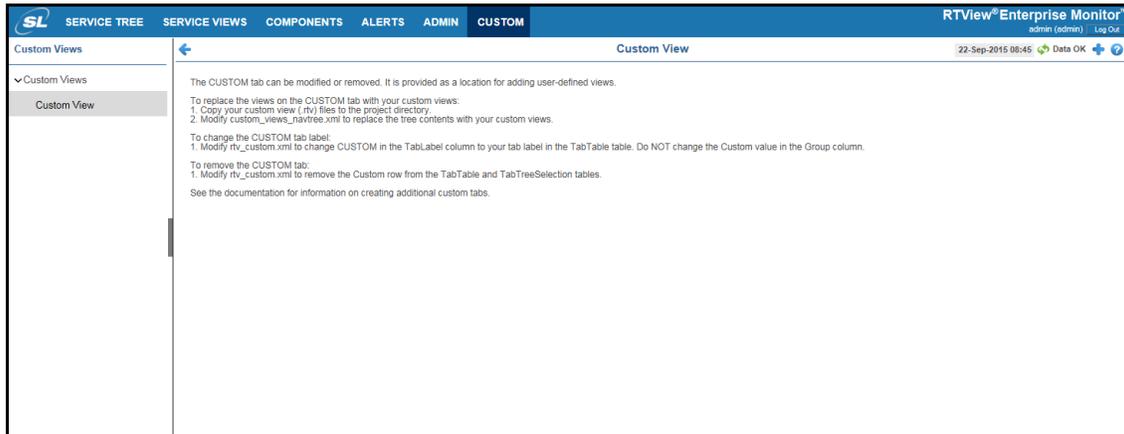
The screenshot shows the 'ADMIN' tab in the RTView Enterprise Monitor. The 'Alert Administration' view is active, displaying a table of alerts. The table has the following columns: Alert, Warning Level, Alarm Level, Duration, Alert Enabled, and Override Count. The 'Alert Enabled' column contains checkboxes, and the 'Override Count' column contains numerical values. Below the table is a 'Settings for Selected Alert' form with the following fields: Name (with a dropdown menu), Description (with a text input and a search icon), Warning Level (with a text input), Alarm Level (with a text input), Duration (Secs.) (with a text input), and Enabled (with a checkbox). A 'Save Settings' button is located at the bottom right of the form.

Alert	Warning Level	Alarm Level	Duration	Alert Enabled	Override Count
AcwInstanceCpuHigh	50	75	30	<input checked="" type="checkbox"/>	0
AcwInstanceDiskReadBytesHigh	100000	200000	30	<input checked="" type="checkbox"/>	0
AcwInstanceDiskReadOpsHigh	100	200	30	<input checked="" type="checkbox"/>	0
AcwInstanceDiskWriteBytesHigh	100000	200000	30	<input checked="" type="checkbox"/>	0
AcwInstanceDiskWriteOpsHigh	100	200	30	<input checked="" type="checkbox"/>	0
AcwInstanceNetworkReadBytesHigh	100000	200000	30	<input checked="" type="checkbox"/>	0
AcwInstanceNetworkWriteBytesHigh	100000	200000	30	<input checked="" type="checkbox"/>	0
AmxServiceHitRateHigh	200	400	30	<input checked="" type="checkbox"/>	0
AmxServiceNodeFailRateHigh	200	400	30	<input checked="" type="checkbox"/>	0
AmxServiceNodeHitRateHigh	200	400	30	<input checked="" type="checkbox"/>	0
AmxServiceNodeMovingAvgHitRateHigh	200	400	30	<input checked="" type="checkbox"/>	0
AmxServiceNodeMovingAvgResponseTimeHigh	200	400	30	<input checked="" type="checkbox"/>	0
AmxServiceNodeResponseTimeHigh	200	400	30	<input checked="" type="checkbox"/>	0
AmxServiceResponseTimeHigh	200	400	30	<input checked="" type="checkbox"/>	0
Bw6AppNodeCpuUsedHigh	50	80	30	<input type="checkbox"/>	0
Bw6AppNodeMemUsedHigh	50	80	30	<input type="checkbox"/>	0
Bw6AppProcessCreatedRateHigh	50	80	30	<input checked="" type="checkbox"/>	0
Bw6AppProcessElapsedTimeHigh	100	200	30	<input type="checkbox"/>	0
Bw6AppProcessExecutionTimeHigh	50	80	30	<input type="checkbox"/>	0
Bw6AppProcessFailedRateHigh	50	80	30	<input type="checkbox"/>	0
Bw6ProcessActivityErrorRateHigh	100	200	30	<input type="checkbox"/>	0
Bw6ProcessCreatedRateHigh	50	80	30	<input type="checkbox"/>	0
Bw6ProcessElapsedTimeHigh	100	200	30	<input type="checkbox"/>	0
Bw6ProcessExecutionTimeHigh	50	80	30	<input type="checkbox"/>	0
Bw6ProcessFailedRateHigh	50	80	30	<input type="checkbox"/>	0
Bw6ProcessSuspendRateHigh	50	80	30	<input type="checkbox"/>	0

The **ADMIN** tab can only be accessed by administrators of RTView Enterprise, who can use this tab during installation to set up proper alert settings, to describe logical and service groupings that drive the construction of the Service Tree, and to “monitor the monitor” view of the current health state of RTView Enterprise and how it is currently deployed. This tab provides access to the **Alert Administration**, **CMDB Administration**, and **Architecture Views**. See the following sections for more information:

- **“Administration”**: Displays in this View allow you to set alert thresholds, track alert management, and modify your Service Data Model.
- **“CMDB Administration”**: Use this display to setup, view, or modify your Service Data Model (CMDB), including: adding, renaming, deleting or merging your CMDB hierarchical elements (Owners, Areas, Groups or Services), associating CIs with Services and assigning or modifying CI attributes (such as Criticality).
- **“Architecture”**: Displays in this View show RTView Enterprise system information such as a topological view of your components and their connection state, configuration definitions and mapping, and performance metrics for your Cache Tables and Data Servers.
- **“Property Views”**: Use this display to see how your properties are configured and the values for all connected RTView processes.

## CUSTOM Tab



The **CUSTOM** tab provides a location where you can add your own custom tab and views, and create diagram displays. See the following sections for more information:

- **Modify the CUSTOM Tab**
- **"Diagram Views"**: Provides the Diagram Generator, which enables you to auto-generate a topology view of your system components.

## Fundamental Structure of Displays

To interpret RTView Enterprise displays it is helpful to understand the Service Data Model. The Service Data Model, also referred to as the CMDB, is a database that forms the fundamental structure of all RTView Enterprise displays, and enables data aggregation and filtering.

The Service Data Model has a four level hierarchy which is, from the highest level (Owner) to the lowest level (Service):

- Owner
- Area
- Group
- Service

The Service Data Model maps all the component IDs (CIs) in your RTView Enterprise system to one or more Services (CIs are items being monitored by RTView Enterprise--servers, processes and so forth--anything that can be configured). Each Service is mapped to a Group, each Group to an Area and each Area to an Owner. Displays are organized and populated with data according to this hierarchy. This mapping enables RTView Enterprise to aggregate data for several hundreds of CIs, and allows objects (heatmaps, tables and so forth) to filter data shown according to user selections.

For details about the configuring the Service Data Model, see the Configure Service Data Model section.

## Heatmaps

Heatmaps organize CIs (according to the Service Data Model) into rectangles and use color to highlight the most critical value in each. Heatmaps enable you to view various alert metrics in the same heatmap using drop-down menus. Each Metric has a color gradient bar that maps relative values to colors. In most heatmaps, the rectangle size represents the number of CIs in the rectangle; a larger size is a larger value.

Heatmaps scale color for a given metric according to the following rules and are applied in the following order:

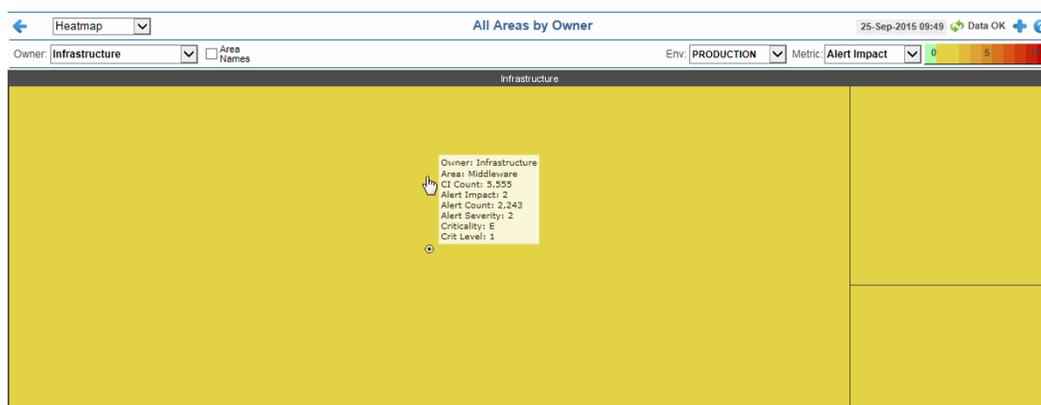
- a) If the metric is associated with an alert, then the color range is scaled from zero to the metric's high Alarm Level threshold, and the color will be red for values near the alarm threshold.
- b) If the metric is not associated with an alert, but the metric is bounded (for example, the **CPU %** utilization value must be in the **0 to 100%** range), then the color is scaled using the user-specified maximum value for the metric.
- c) Otherwise, the metric is autoscaled into a color range from white (minimum) to green (high) using the current highest metric value observed over the monitored entities.

By default, the metric is linearly scaled to an appropriate color. If the **Log** checkbox is checked, then the selected color reflects the logarithm of the current metric value.

Heatmaps include drop-down menus to filter data by Owner, Area, Group, Service, Region and Environment. The filtering options vary among heatmaps.

For example, the **All Management Areas - "Area Heatmap"** (shown in the following figure) illustrates a typical RTView Enterprise heatmap. The heatmap contains a **Metric** drop-down menu with options to show **Alert Impact**, **Alert Severity**, **Alert Count** and **Criticality** (menu options vary according to the data populating the heatmap). **Alert Impact** is selected and its corresponding color gradient bar  is shown. Each rectangle represents all CIs in an Area. The red rectangle in the heatmap indicates that one or more CIs in that Area currently has an alert in an alarm state. The yellow rectangles in the heatmap indicate that one or more CIs in those Areas currently have an alert in a warning state. A green rectangle would indicate that no alert is in a warning or alarm state in an Area.

Continuing with our example, there are two filtering options. You can choose to show all Owners or a single Owner, and all Environments or a single Environment. Each rectangle represents an Area. The rectangle size represents the number of CIs in the rectangle; a larger size is a larger value. Use the check-boxes  to include or exclude labels in the heatmap. Move your mouse over a rectangle to see additional information. The following figure illustrates the mouse-over feature in which we see all the **Metric** drop-down values.



In most heatmaps, you can also drill-down to more detail by clicking a rectangle in the heatmap. Or, click Open New Window  and then drill-down. The drill-down opens a display that contains relevant and more detailed data.

**Filter By:**

The following filtering options are typically included:

**Owner:** Choose an Owner to see metrics for Areas associated with that Owner.

**Area:** Choose an Area to see metrics for Groups associated with that Area and Owner.

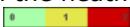
**Group:** Choose a Group to see metrics for Services associated with that Group, Area and Owner.

**Service:** Choose a Service to see metrics for Environments associated with that Service, Group, Area and Owner.

**Env:** Choose an Environment to see metrics for Environments associated with that Service, Group, Area and Owner.

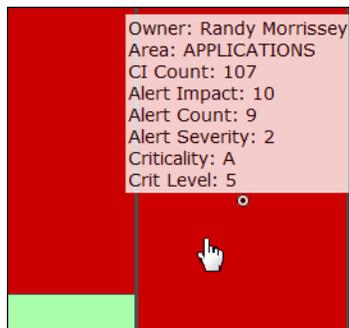
**Metric:**

Choose the type of metric to show in the heatmap. Each metric has its own gradient bar that maps relative values to colors:

<b>Alert Impact</b>	The product of the maximum Alert Severity of alerts in the heatmap rectangle multiplied by the maximum Criticality of alerts in the heatmap rectangle. Values range from <b>0</b> - <b>10</b> , as indicated in the color gradient  bar, where <b>10</b> is the highest Alert Impact.
<b>Alert Severity</b>	<p>The maximum level of alerts in the heatmap rectangle. Values range from <b>0</b> - <b>2</b>, as indicated in the color gradient  bar, where <b>2</b> is the highest Alert Severity.</p> <ul style="list-style-type: none"> <li> Red indicates that one or more metrics have reached their alarm threshold. Metrics that have exceeded their specified ALARM LEVEL threshold have an Alert Severity value of <b>2</b>.</li> <li> Yellow indicates that one or more metrics have reached their alarm threshold. Metrics that have exceeded their specified WARNING LEVEL threshold have an Alert Severity value of <b>1</b>.</li> <li> Green indicates that no metrics have reached their alert thresholds. Metrics that have not exceeded their specified thresholds have an Alert Severity value of <b>0</b>.</li> </ul>
<b>Alert Count</b>	The total number of critical and warning alerts in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.
<b>Criticality</b>	<p>The maximum level of Criticality (rank of importance) in the heatmap rectangle. Values range from <b>1</b> to <b>5</b>, as indicated in the color gradient  bar, where <b>5</b> is the highest Criticality.</p> <p>Criticality is specified in the Service Data Model (CMDB) by your administrator. Criticality values are listed in the <b>Component Views</b> - "<a href="#">CI / Service Table</a>" display, which range from <b>A</b> to <b>E</b>, where <b>A</b> is the highest Criticality (level <b>5</b> maps to a Criticality of <b>A</b> and level <b>1</b> maps to a Criticality of <b>E</b> with equally spaced intermediate values).</p>

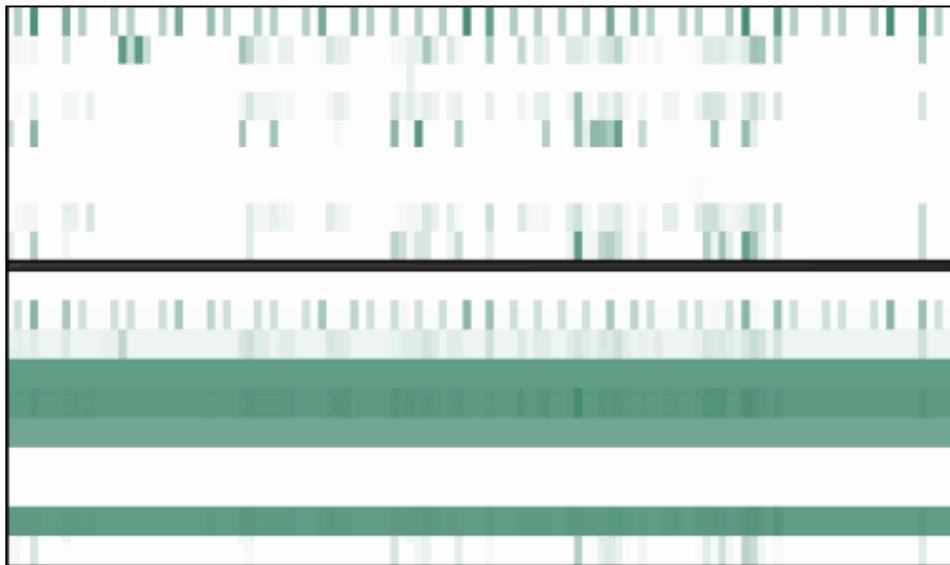
## Mouse-over

The mouse-over functionality provides additional detailed data in an over imposed pop-up window when you mouse-over a heatmap. The following figure illustrates mouse-over functionality in a heatmap object. In this example, when you mouse-over a host, details are shown such as **CI Count**, **Alert Impact**, **Alert Severity**, and **Criticality**.



## History Heatmaps

History heatmaps, such as the following cache heatmap, show you utilization trends, over time, for an entire Oracle Coherence cluster.



Each row represents a cache. Each column represents a time period. A darker color indicates heavier usage, a lighter color indicates lighter usage. At a glance, you can quickly analyze load distribution, check for bottlenecks and identify caches with high usage. You can also answer questions such as, Is the cluster using what I expect? Is the cluster using it in a uniform scale? If there is an issue, you can mouse-over the heatmap to see when the issue started, what behavior preceded it, and the name of the resource.

Additionally, because data updates for all the elements in your cluster share the same time-stamp, you can see utilization spikes in the cluster, such as in trend graphs or heatmaps, and immediately address performance issues. Other monitoring systems cannot gather enough simultaneous data points for displaying spikes.

## Tables

Tables contain the same data that is shown in the heatmap in the same View, and additional data not included the heatmap.

All Properties			
Order	File Name	Property Name	Property Value
2		sl.rtvview.013420134501337013390134401302	01350013480135101335
2	Sort Ascending	sl.rtvview.cache.config	adaptris_cache_source.rtv.\$adap
2	Sort Descending	sl.rtvview.jmx.jmxconn	demoAdapter - - URL:service:jmx
2		rtvapm_package	adaptrismon
2	Columns	rtv_proctag	ADAPTRISMON
2		rtv_title	ADAPTRISMON Package
2	Filter	sl.rtvview.sql.sqldb	ALERTDEFS sa - jdbc:hsqldb:hs
2		sl.rtvview.cp	C:/rtvdemos/mysql-connector-jav
2	Settings	historian.sl.rtvview.historian.driver	com.mysql.jdbc.Driver
213	rtvview	historian.sl.rtvview.historian.url	jdbc:mysql://192.168.200.42:3306
215	rtvview	historian.sl.rtvview.historian.password	my-secret-pw

Tables support advanced HTML interactive features such as sorting on multiple columns, filtering on multiple columns, column resizing, column reordering, and hiding columns. Many of these features are accessed from the column menu, shown in the screen shot above, which you open by clicking on the menu icon in a column's header.

Some tables in the **Components** tab gray out rows when they're in an expired state. A row is expired when data has not been received within the time specified in the solution package that is hosting the data.

Also see:

- ["Multiple Column Sorting"](#)
- ["Column Visibility"](#)
- ["Column Filtering"](#)
- ["Column Locking"](#)
- ["Column Reordering"](#)
- ["Saving Settings"](#)
- ["Row Paging"](#)
- ["Row Color Code"](#)
- ["Row Keyboard Selection"](#)

### Multiple Column Sorting

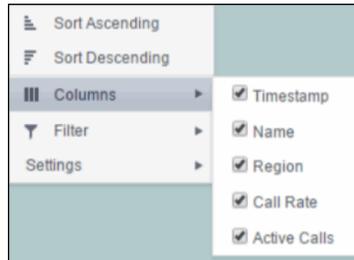
Click on a column header to sort the table by that column. On the first click, the column is sorted in ascending order (smallest value at the top), on the second click the sort is in descending order, and on the third click, the column is returned to its original unsorted state. A sort on a string column is case-insensitive.

To sort multiple columns, click on the column header for each column you want to sort. The sorting is performed in the order that the column headers were clicked. Multiple column sorting is a very useful feature, but can also cause confusion if you intend to sort on a single column, but forget to "unsort" any previously selected sort columns first. You should check for the up/down sort icon in other column headers if a sort gives unexpected results.

The grid's row selection is cleared if the sort is changed or if columns are resized or reordered. Column sorting is reflected in an export to HTML and Excel.

## Column Visibility

You can hide or show columns in the table by clicking on any column's menu icon, and choosing **Columns** from the menu. This opens a submenu with a check box for each column that toggles the visibility of the column. All columns in the data table appear in the Columns menu, even those that are initially hidden.



The leftmost column (the row header column) cannot be hidden.

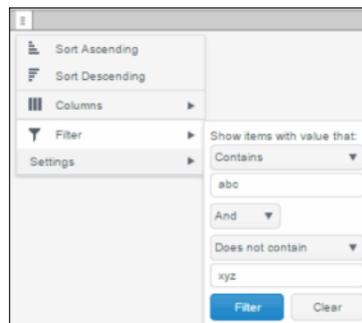
Column visibility changes are NOT reflected in an export to HTML and Excel.

## Column Filtering

You can create a filter on any column. If filters are created on multiple columns, then only the rows that pass all of the filters are displayed. That is, if there are multiple filters they are logically "ANDed" together to produce the final result.

The background of a column's menu icon changes to white to indicate that a filter is defined on that column. This is intended to remind you which columns are filtered.

You can configure a filter on any column by clicking on the column's menu icon and choosing **Filter** from the menu. This opens the **Column Filter** dialog:



Options in the **Column Filter** dialog vary according to the data type of the selected column:

- String columns:** You can enter a filter string such as "abc" and, from the dropdown list, select the operator (equal to, not equal to, starts with, contains, etc) to be used when comparing the filter string to each string in the column. All of the filter comparisons on strings are case-insensitive. You can optionally enter a second filter string (e.g. "xyz") and specify if an AND or OR combination should be used to combine the first and second filter results on the column.

- **Numeric columns:** You can enter numeric filter values and select arithmetic comparison operators, ( $=$ ,  $\neq$ ,  $>$ ,  $\geq$ ,  $<$ ,  $\leq$ ). You can optionally enter a second filter value and comparison operator, and specify if an AND or OR combination should be used to combine the first and second filter results.
- **Boolean columns:** You simply select whether matching items should be true or false.

The numeric and boolean filter dialogs are shown below.

- **Date columns:** You can select a date and time and choose whether matching items should have a timestamp that is the same as, before, or after the filter time. The date is selected by clicking on the calendar icon and picking a date from a calendar dialog. The time is selected by clicking on the time icon and picking a time from a dropdown list:

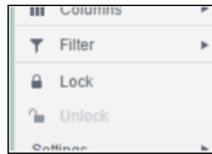
Alternatively, a date and time can be typed into the edit box. The strings shown in a date column are formatted by the Display Server using its time zone. But if a filter is specified on a date column, the date and time for the filter are computed using the client system's time zone. This can be confusing if the Display Server and client are in different time zones.

Data updates to the grid are suspended while the filter menu is opened. The updates are applied when the menu is closed.

Column filtering is reflected in an export to HTML and Excel.

## Column Locking

The leftmost column is "locked" in position, meaning that it does not scroll horizontally with the other columns in the table. If the row header is enabled, then two items labeled **Lock** and **Unlock** appear in the column menu. These can be used to add or remove additional columns from the non-scrolling row header area.



If the row header is enabled, at least one column must remain locked.

Column locking is NOT reflected in an export to HTML and Excel.

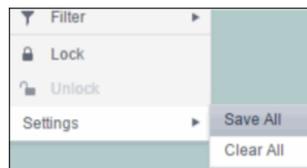
## Column Reordering

You can reorder the grid columns by dragging and dropping a column's header into another position. Dragging a column into or out of the row header area (the leftmost columns) is equivalent to locking or unlocking the column.

Column reordering is NOT reflected in an export to HTML and Excel.

## Saving Settings

You can permanently save all of the custom settings made to the grid, including filtering, sorting, column size (width), column order, column visibility, and column locking. This is done by opening any column menu, clicking **Settings**, and then clicking **Save All**:



The grid's settings are written as an item in the browser's local storage. The item's value is a string containing the grid's settings. The item uses a unique key comprised of the URL path name, the display name, and the table's RTView object name. If the Thin Client's login feature is enabled, the key will also include the username and role, so different settings can be saved for each user and role for a grid on any given display, in the same browser and host.

If you save the grid settings and navigate away from the display or close the browser, then the next time you return to the display in the same browser the settings are retrieved from the browser's local storage and applied to the grid. The browser's local storage items are persistent, so the grid settings are preserved if the browser is closed and reopened or if the host system is restarted.

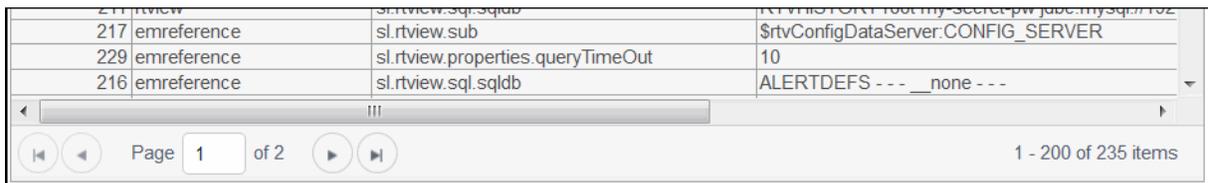
Note that each browser has its own local storage on each host. The local storage items are not shared between browsers on the same host or on different hosts. So, if a user logs in as Joe with **role = admin**, in Internet Explorer on host H1, and saves grid settings for display X, then those grid settings are restored each time a user logs in as Joe, role admin, on host H1 and opens display X in Internet Explorer. But if all the same is true except that the browser is Chrome, then the settings saved in Internet Explorer are not applied. Or if the user is Joe and role is admin and the browser is IE and the display is X, but the host system is H2 not H1, then the grid settings saved on H1 are not applied.

### Revert Table Settings

You can delete the grid's item from local storage by clicking **Settings> Clear All** in any column menu. This permanently deletes the saved settings for the grid and returns the grid to the state defined in the display file.

### Row Paging

If the data table contains more than one 200 rows, page controls appear at the bottom of the grid.



### Row Color Code

Table rows sometimes use color to indicate the current most critical alert state for all CIs associated with the row. In this example, the **Severity Level** column is sorted in descending order (from high to low values).

JVM	localhostGLASSFISH_SERVER_8	10	10
JVM	localhostMYDEMO_DATASERVER	8	8
JVM	localhostMYDEMO_DISPLAYSERVER	8	8
JVM	slidemos.com213415_RTVD8	10	10
JVM	localhostIWM-DB-1	5	5
WAS	SLHOST12Node01Cell SLHOST12Node01.server1	5	5
JVM	localhostRTVMGR_DATABASE	5	5
JVM	localhostRTVMGR_DATASERVER	0	0
JVM	localhostWLM_DATABASE	0	0
EMS	tcp:SLHOST10.7021	0	0
EMS	tcp:SLHOST10.7020	0	0
WLS	TestDomain.ManagedServer2	0	0

The yellow row color indicates that one or more alerts exceeded their warning threshold for one or more CIs associated with the Service. The red row color indicates that one or more alerts exceeded their critical threshold for the CI associated with the Service (in this case there

is a single CI). To summarize:

**Row Color Code:**

Tables with colored rows indicate the following:

- Red indicates that one or more alerts exceeded their ALARM LEVEL threshold in the table row.
- Yellow indicates that one or more alerts exceeded their WARNING LEVEL threshold in the table row.
- Green indicates that no alerts exceeded their WARNING or ALARM LEVEL threshold in the table row.

## Row Keyboard Selection

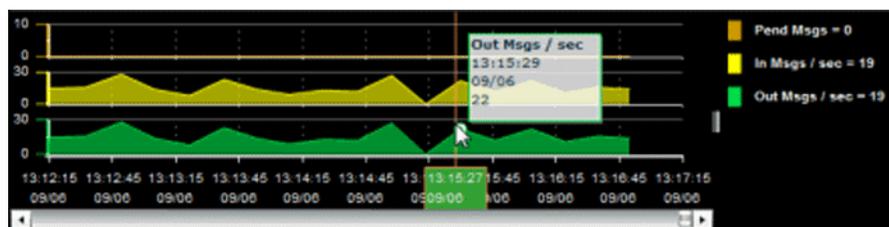
You can use the mouse to select a row and use the arrow keys to change the focus (highlighted) row, but to select the focus row, you must then press the space bar.

8	C:\rtvdemos\rtvapm\common\conf\rtvapm	sl.rtvview.sql.dbretry
9	C:\rtvdemos\rtvapm\common\conf\rtvapm	sl.rtvview.global
10	C:\rtvdemos\rtvapm\common\conf\rtvapm	sl.rtvview.global
11	C:\rtvdemos\rtvapm\common\conf\rtvapm	sl.rtvview.xml.xmlsource
12	C:\rtvdemos\rtvapm\common\conf\rtvapm	sl.rtvview.jmx.jmxconn
13	C:\rtvdemos\rtvapm\common\conf\rtvapm	sl.rtvview.dsenable

## Trend Graphs

Trend graphs enable you to view and compare various important metrics over time, such as server memory utilization, server throughput, the number of clients being served by the server, or the total amount of data sent to clients. You can use trend graphs to assess utilization and performance trends.

For example, the following figure illustrates a typical trend graph. In this example, metrics for **Pending Messages**, **Incoming Messages** and **Outgoing Messages** are traced.



By default, the time range end point is the current time. To change the time range for the trend graph click Open Calendar , choose the date and time, then click **OK**. Or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, **Apr 26, 2012 5:01 PM**. Click **Apply**. Use the Navigation Arrows   to move forward or backward one time period (the time period selected from the **Time Range** drop-down menu). Click **Restore to Now** to reset the time range end point to the current time.

## Mouse-over

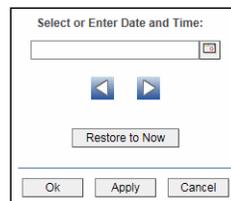
The mouse-over functionality provides additional detailed data in an over imposed pop-up window when you mouse-over trend graphs. The above figure illustrates mouse-over functionality. In this example, when you mouse-over a single dot, or data point, in the **Out Msgs / sec** trend graph, a pop-up window shows data for that data point. In this case, the X-axis value is **13:15:29 hours on September 6th**, and the Y-axis value is **22 Outbound messages per second**.

## Log Scale

Typically, trend graphs provide the Log Scale option. Log Scale enables you to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.

## Time Range

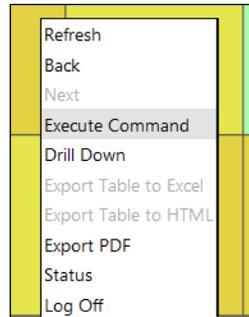
Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. By default, the time range end point is the current time. To enter a specific time range, click the associated ellipsis button .



To change the time range click the Open Calendar button , choose the date and time, then click **OK**. Or, enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss** (for example, Aug 21, 2011 12:24 PM) and click **Apply**. Use the Navigation Arrows   to move forward or backward one time period (the time period selected from the Time Range drop-down menu). Click **Restore to Now** to reset the time range end point to the current time.

## Popup Menu

Typically, you can right-click on displays to open a popup menu. By default, options include **Refresh**, **Back**, **Next**, **Execute Command**, **Drill Down**, **Export Table to Excel**, **Export Table to HTML**, **Export PDF**, **Status** and **Log Off**. The following figure illustrates the popup menu in a heatmap.

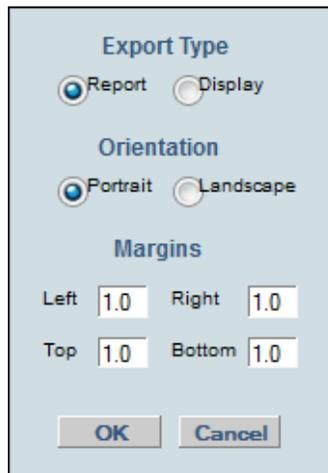


## Export PDF Report

You can quickly export reports for displays, or for tables and grid objects in a display, to a PDF file.

### To generate a report for a display:

Right-click on the display and select **Export PDF**. The **Export to PDF** dialog opens.

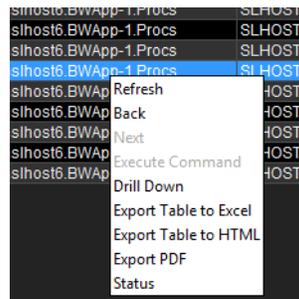


Set the margins and choose the **Export Type**:

- Report:** Generates an image of the display on the first page, followed by at least one page for each table or object grid in the display. As many pages as are necessary to show all the data in each table or object grid are included in the report. This enables you to view all data in a table or object grid that you otherwise must use a scrollbar to see. If there are no tables or object grids in your display, you only get a image of the display.
- Display:** Generates an image of the display in PDF format. Choose the page orientation (**Portrait** or **Landscape**), set the page margins and click **OK**. The report opens in a new window.

### To generate a report for a table or grid object in a display:

Right-click on the table or grid object and choose **Export PDF**, **Export Table to Excel** or **Export Table to HTML**.



## Title Bar

Displays share the same top layer in the title bar, as shown below.



The following table describes the functionality in the display title bar.



Opens the previously open display.



Opens the display that is up one level.



Navigates to a display that is most commonly accessed from the current display. The target display differs among displays.



Navigates to displays that are most commonly accessed from the current display. The drop-down menu options differ among displays.



Opens the Alerts Table display in a new window.



The current date and time. If the time is incorrect, this might indicate that RTView stopped running. When the date and time is correct and the **Data OK** indicator is green, this is a strong indication that the platform is receiving current and valid data.



The data connection state. Red indicates the data source is disconnected (for example, if the Data Server is not receiving data, or if the Display Server does not receive data from the Data Server, this will be red). Green indicates the data source is connected. When the date and time is correct and the **Data OK** indicator is green, this is a strong indication that the platform is receiving current and valid data.



Opens an instance of the same display in a new window. Each window operates independently, allowing you to switch views, navigate to other displays in RTView EM, and compare server performance data. For illustration, see **Multiple Windows**.



Opens the online help page for the current display.

6,047

The number of items (for example, CIs or Areas) in the display.

Area Count: 9

## Multiple Windows

The following illustrates the usage of the Open New Window  to open multiple windows.

The screenshot shows three overlapping windows in the RTView EM interface:

- Top Window:** "All Areas by Owner" showing a heatmap view of service areas.
- Middle Window:** "All Services Status History" displaying a list of services with their status (Quality, Severity, Alerts, Criticality, Impact) and Region. The "CI Type Filter" is set to "All CI Types".
- Bottom Window:** "Single Service Summary" for "IRIS-SCAN-DFW". It shows service details and a table of CI (Configuration Item) data.

CIType	CIName	Quality	Severity	Alerts	Criticality	Impact	Region
EMS-QUEUE	tcp://MIRIS1034.7222:SCAN-QUEUE	●	●	0	C	0	AMER
EMS-QUEUE	tcp://MIRIS1034.7222:WORK-QUEUE	●	●	0	C	0	AMER
EMS-QUEUE	tcp://MIRIS1075.7222:REPORT-QUEUE	●	●	0	C	0	AMER
EMS-SERVER	tcp://MIRIS1034.7222	●	●	0	C	0	AMER
JVM	localhost:VMIRIS1034-SCAN-JVM-DFW	●	●	0	C	0	AMER
VMWARE-HOST	vSphereS:esxi-1.south	●	●	0	C	0	AMER
VMWARE-VM	vSphereS:VMIRIS1034	●	●	0	C	0	AMER
WLS	Domain-SOUTH:WLS-SERVER-DFW	●	●	0	C	0	AMER
WLS-APP	Domain-SOUTH:WLS-SERVER-DFW:control	●	●	0	C	0	AMER
WLS-APP	Domain-SOUTH:WLS-SERVER-DFW:scanner	●	●	0	C	0	AMER

---

## RTViewCentral Displays

This section describes the Views and displays that come with RTView Enterprise and reside on RTViewCentral. This section describes the following Views:

- ["All Management Areas"](#)
- ["Multi Area Service Views"](#)
- ["Single Area Service Views"](#)
- ["Service Summary Views"](#)
- ["Key Metrics Views"](#)
- ["Component Views"](#)
- ["Metric Explorer"](#)
- ["Alert Views"](#)
- ["Administration"](#)
- ["CMDB Administration"](#)
- ["Architecture"](#)
- ["Property Views"](#)
- ["Diagram Views"](#)
- ["Common/Alerts - HTML"](#)
- ["Common/System - HTML"](#)

For details about add-on displays, see:

- ["RTView DataServer for IBM"](#)
- ["RTView DataServer for Infrastructure"](#)
- ["RTView DataServer for Kafka"](#)
- ["RTView DataServer for Oracle"](#)
- ["RTView DataServer for Solace"](#)
- ["RTView DataServer for TIBCO"](#)

### All Management Areas

These displays present the highest-level summary views of alert states for your entire system. Aggregated data is organized by Owners and shows all Areas, while highlighting the most critical alert states using color. Data can be filtered by Owner, Area, Environment and alert Metric. Data is filtered by the \$rtvOwnerMask and \$rtvAreaMask values for the logged in user. For details, refer to the *RTView Enterprise Configuration Guide*.

Use these displays to monitor critical alerts anywhere in your system, and investigate those alerts in lower-level displays. Because these displays immediately show you any critical alert in your system, users typically keep one of these displays open for quick monitoring. Click an Area in the display to drill-down and view the selected Area in the **Multi Area Service Views** displays.

The **All Management Areas** displays present the same aggregated data in tabular and heatmap formats. Displays in this View are:

- **"Area Heatmap"**: Heatmap of the most critical alerts for all Areas of your system, with the option to filter by Owner, Environment and alert Metric.
- **"Area Table"**: Table of data shown in the **All Management Areas - "Area Heatmap"** with the option to filter by Owner and Environment.

## Area Heatmap

View the most critical alert state for all monitored instances throughout your system. Consider keeping this display open to monitor conditions in your system. The heatmap organizes monitored instances by one or all Owners for all Areas, and uses color to show the most critical alert state in each. Each rectangle in the heatmap represents a management Area (for example, Applications, Demo Systems and so forth), which are also grouped by Owner. The rectangle size represents the number of CIs in the rectangle; a larger size is a larger value.

Use the available drop-down menus or right-click to filter data shown in the display. Use the check-boxes  to include or exclude labels in the heatmap. Move your mouse over a rectangle to see additional information. By default, this display shows all Owners, all Environments and the Alert Impact.

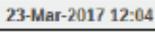
Drill-down and investigate by clicking a rectangle in the heatmap to view details for the selected Area in the display that was last selected under **Multi Area Service Views**. For example, if the last selected display under **Multi Area Service Views** was **"Group / Service Table"**, then clicking an Area in the heatmap results in displaying details in the **Group / Service Table** display.



### Title Bar (possible features are):

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu ▾, Table open commonly accessed displays.

 Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

**Filter By:**

The following filtering options are typically included:

**Owner:** Choose an Owner to see metrics for Areas associated with that Owner.

**Area:** Choose an Area to see metrics for Groups associated with that Area and Owner.

**Group:** Choose a Group to see metrics for Services associated with that Group, Area and Owner.

**Service:** Choose a Service to see metrics for Environments associated with that Service, Group, Area and Owner.

**Env:** Choose an Environment to see metrics for Environments associated with that Service, Group, Area and Owner.

**Metric:**

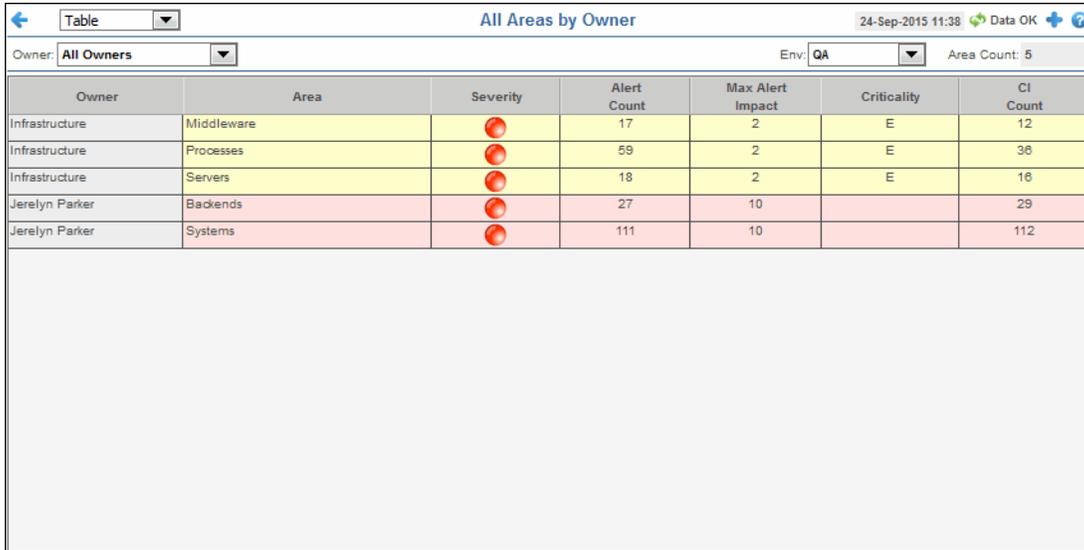
Choose the type of metric to show in the heatmap. Each metric has its own gradient bar that maps relative values to colors:

<b>Alert Impact</b>	The product of the maximum Alert Severity of alerts in the heatmap rectangle multiplied by the maximum Criticality of alerts in the heatmap rectangle. Values range from <b>0</b> - <b>10</b> , as indicated in the color gradient  bar, where <b>10</b> is the highest Alert Impact.
<b>Alert Severity</b>	The maximum level of alerts in the heatmap rectangle. Values range from <b>0</b> - <b>2</b> , as indicated in the color gradient  bar, where <b>2</b> is the highest Alert Severity. <ul style="list-style-type: none"> <li><span style="color: red;">●</span> Red indicates that one or more metrics have reached their alarm threshold. Metrics that have exceeded their specified ALARM LEVEL threshold have an Alert Severity value of <b>2</b>.</li> <li><span style="color: yellow;">●</span> Yellow indicates that one or more metrics have reached their alarm threshold. Metrics that have exceeded their specified WARNING LEVEL threshold have an Alert Severity value of <b>1</b>.</li> <li><span style="color: green;">●</span> Green indicates that no metrics have reached their alert thresholds. Metrics that have not exceeded their specified thresholds have an Alert Severity value of <b>0</b>.</li> </ul>
<b>Alert Count</b>	The total number of critical and warning alerts in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.
<b>Criticality</b>	The maximum level of Criticality (rank of importance) in the heatmap rectangle. Values range from <b>1</b> to <b>5</b> , as indicated in the color gradient  bar, where <b>5</b> is the highest Criticality. <p>Criticality is specified in the Service Data Model (CMDB) by your administrator. Criticality values are listed in the <b>Component Views</b> - "<a href="#">CI / Service Table</a>" display, which range from <b>A</b> to <b>E</b>, where <b>A</b> is the highest Criticality (level <b>5</b> maps to a Criticality of <b>A</b> and level <b>1</b> maps to a Criticality of <b>E</b> with equally spaced intermediate values).</p>

**Area Table**

View data from the **All Management Areas** - "[Area Heatmap](#)" in a tabular format: all alert states (alert Impact, Severity, Count, Criticality and CI Count) for all Areas, Owners and Environments. Each row in the table is a different Area (for example, **Applications**, **Demo Systems** and so forth). Use this display to check the status of your systems by Area, Owner and Environment, and to compare detailed metrics across all Areas in your organization.

Use the available drop-down menus or right-click to filter data shown in the display. Click Sort  to order column data. Drill-down and investigate by clicking a row in the table to view details for the selected Area in the display that was last selected under **Multi Area Service Views**. For example, if the last selected display under **Multi Area Service Views** was "Group / Service Table", then clicking an Area in the heatmap results in displaying details in the **Group/Service Table** display.



Owner	Area	Severity	Alert Count	Max Alert Impact	Criticality	CI Count
Infrastructure	Middleware		17	2	E	12
Infrastructure	Processes		59	2	E	38
Infrastructure	Servers		18	2	E	16
Jerelyn Parker	Backends		27	10		29
Jerelyn Parker	Systems		111	10		112

#### Title Bar (possible features are):

-   Open the previous and upper display.
-  Open an instance of this display in a new window.
-  Open the online help page for this display.
- open commonly accessed displays.

 Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

#### Row Color Code:

Tables with colored rows indicate the following:

-  Red indicates that one or more alerts exceeded their ALARM LEVEL threshold in the table row.
-  Yellow indicates that one or more alerts exceeded their WARNING LEVEL threshold in the table row.
-  Green indicates that no alerts exceeded their WARNING or ALARM LEVEL threshold in the table row.

**Filter By:**

The following filtering options are typically included:

**Owner:** Choose an Owner to see metrics for Areas associated with that Owner.

**Area:** Choose an Area to see metrics for Groups associated with that Area and Owner.

**Group:** Choose a Group to see metrics for Services associated with that Group, Area and Owner.

**Service:** Choose a Service to see metrics for Environments associated with that Service, Group, Area and Owner.

**Env:** Choose an Environment to see metrics for Environments associated with that Service, Group, Area and Owner.

**Fields and Data**

This display includes:

**Area Count** The current number of Areas shown in the table.

**(Table)**

Each row in the table is a different Area.

<b>Owner</b>	The name of the person or Group the Area is designated to.
<b>Area</b>	The name of the Area where the alert data originated.
<b>Severity</b>	The maximum level of alerts in the Area. Values range from <b>0</b> to <b>2</b> , where <b>2</b> is the greatest Severity: <ul style="list-style-type: none"> <li><span style="color: red;">●</span> One or more alerts exceeded their ALARM LEVEL threshold in the Area.</li> <li><span style="color: yellow;">●</span> One or more alerts exceeded their WARNING LEVEL threshold in the Area.</li> <li><span style="color: green;">●</span> No alert thresholds have been exceeded in the Area.</li> </ul>
<b>Criticality</b>	The Criticality (rank of importance) specified in the Service Data Model (CMDB) by your administrator. Criticality values are listed in the Component Views / CI Service Table display, which range from A to E, where A is the highest Criticality. This value is used to determine the value for Alert Impact.
<b>Max Alert Impact</b>	The highest value that Alert Impact has had for the Area.
<b>Alert Count</b>	The total number of critical and warning alerts for the Area.
<b>CI Count</b>	The total number of configurable items associated with the Area.

## Multi Area Service Views

These displays present aggregated data of alert states for all Services for all Areas. Data can be filtered by Area, Group, Environment, and alert Metric. Data is filtered by the \$rtvOwnerMask, \$rtvAreaMask and \$rtvGroupMask values for the logged in user. For details, refer to the *RTView Enterprise Configuration Guide*.

Use these displays, for example, to isolate the Area and Environment in which a critical alert is occurring. If you see a critical alert, get information by comparing alert metrics (such as how many other items are potentially affected).

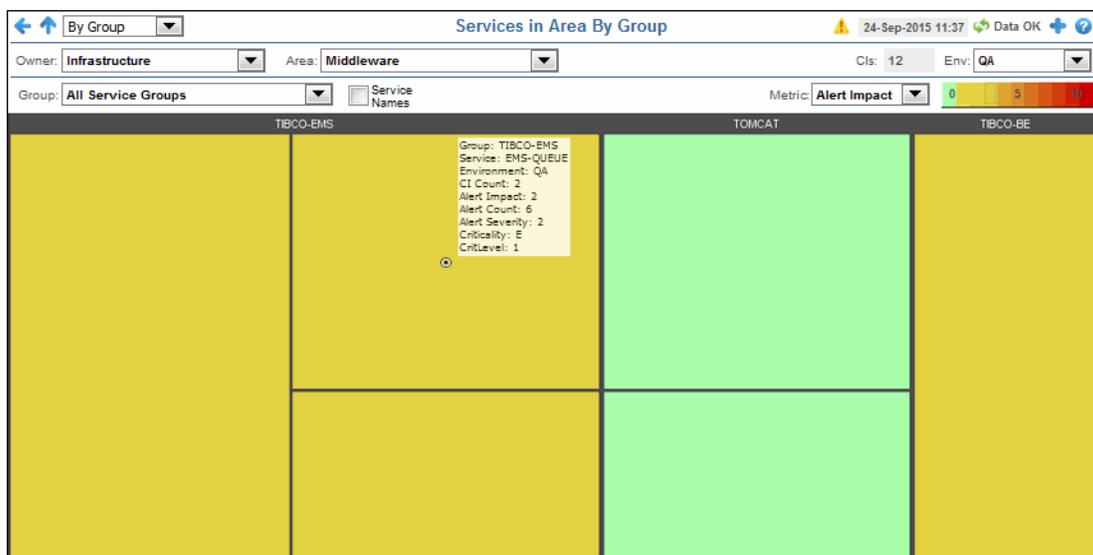
These displays drill-down to the **Service Summary Views** - “[Service By CI Type](#)” display. The **Multi Area Service Views** displays present data in tabular and heatmap formats. Displays in this View are:

- “[Group/Service Heatmap](#)”: Heatmap of alert states for Services by Area, with the option to filter by Area, Group, Environment and alert Metric, and the option to show Group and Service Names.
- “[Group/Region Heatmap](#)”: Heatmap as described for the **Group / Service Heatmap** (above), with the option to filter by Region and no option to show Service Names.
- “[Group / Service Table](#)”: Table of **Group/Service Heatmap** data.
- “[Services CI Type Summary](#)”: Table that shows the health state of Services per CI Type.
- “[Services History Heatmap](#)”: Heatmap of alert states, over time, for Services in a selected Area, with the option to filter by Group, Environment and alert Metric.

## Group/Service Heatmap

View heatmap of alert states for Services in one or all Areas, filter by Group or Environment, and optionally show Service Names. The heatmap organizes Services by one or all Areas. Each rectangle in the heatmap represents a Service (for example, Applications, Demo Systems and so forth), which are grouped by Area. The rectangle size represents the number of CIs in the Service; a larger size is a larger value.

Use the available drop-down menus or right-click to filter data shown in the display. Use the check-boxes  to include or exclude labels in the heatmap. Move your mouse over a rectangle to see additional information. Drill-down and investigate by clicking a rectangle in the heatmap to view details in the last display that was viewed under either the **Service Summary Views** or **Key Metrics Views**. For example, if the last selected display was the “[Service Summary](#)” display under “[Service Summary Views](#)” and you clicked on a rectangle in the **Group / Service Heatmap**, the details would display in the **Service Summary** display. If the last selected display was the “[Service KM Table](#)” display under “[Key Metrics Views](#)”, then clicking a rectangle in the **Group / Service Heatmap** displays the details in the **Service KM Table**.





**Note:** The "Up" Arrow (↑) opens the most recently viewed display under "All Management Areas". For example, if the last viewed display under **All Management Areas** was **Area Table**, then clicking ↑ opens the "Area Table" display.

### Filter By:

The following filtering options are typically included:

**Owner:** Choose an Owner to see metrics for Areas associated with that Owner.

**Area:** Choose an Area to see metrics for Groups associated with that Area and Owner.

**Group:** Choose a Group to see metrics for Services associated with that Group, Area and Owner.

**Service:** Choose a Service to see metrics for Environments associated with that Service, Group, Area and Owner.

**Env:** Choose an Environment to see metrics for Environments associated with that Service, Group, Area and Owner.

### Metric:

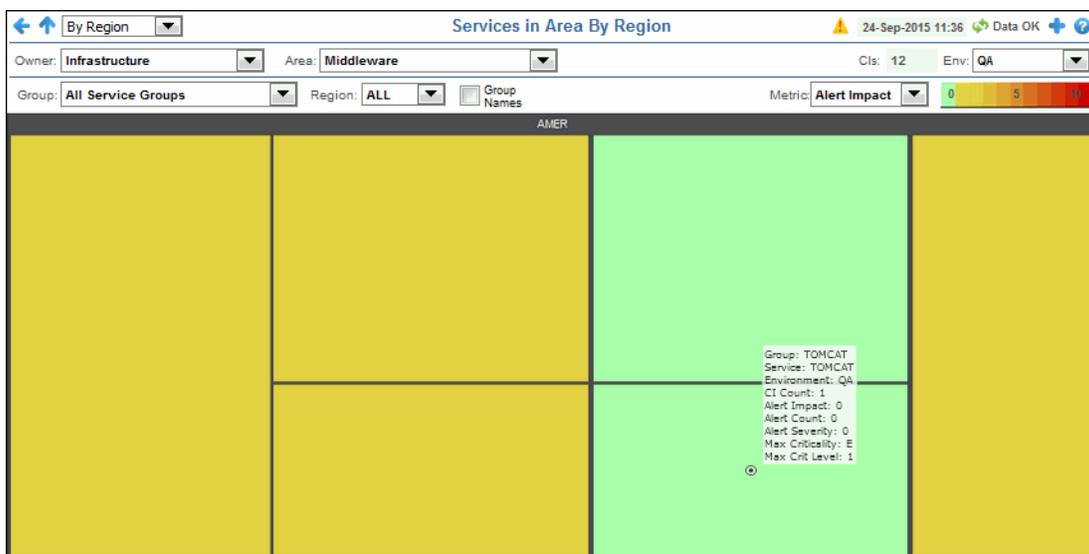
Choose the type of metric to show in the heatmap. Each metric has its own gradient bar that maps relative values to colors:

- |                       |  |
|-----------------------|--|
| <b>Alert Impact</b>   | The product of the maximum Alert Severity of alerts in the heatmap rectangle multiplied by the maximum Criticality of alerts in the heatmap rectangle. Values range from <b>0</b> - <b>10</b> , as indicated in the color gradient  bar, where <b>10</b> is the highest Alert Impact.   |
| <b>Alert Severity</b> | The maximum level of alerts in the heatmap rectangle. Values range from <b>0</b> - <b>2</b> , as indicated in the color gradient  bar, where <b>2</b> is the highest Alert Severity.<br><ul style="list-style-type: none"> <li><span style="color: red;">●</span> Red indicates that one or more metrics have reached their alarm threshold. Metrics that have exceeded their specified ALARM LEVEL threshold have an Alert Severity value of <b>2</b>.</li> <li><span style="color: yellow;">●</span> Yellow indicates that one or more metrics have reached their alarm threshold. Metrics that have exceeded their specified WARNING LEVEL threshold have an Alert Severity value of <b>1</b>.</li> <li><span style="color: green;">●</span> Green indicates that no metrics have reached their alert thresholds. Metrics that have not exceeded their specified thresholds have an Alert Severity value of <b>0</b>.</li> </ul> |
| <b>Alert Count</b>    | The total number of critical and warning alerts in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.   |
| <b>Criticality</b>    | The maximum level of Criticality (rank of importance) in the heatmap rectangle. Values range from <b>1</b> to <b>5</b> , as indicated in the color gradient  bar, where <b>5</b> is the highest Criticality.<br><br>Criticality is specified in the Service Data Model (CMDB) by your administrator. Criticality values are listed in the <b>Component Views</b> - "CI / Service Table" display, which range from <b>A</b> to <b>E</b> , where <b>A</b> is the highest Criticality (level <b>5</b> maps to a Criticality of <b>A</b> and level <b>1</b> maps to a Criticality of <b>E</b> with equally spaced intermediate values).   |

## Group/Region Heatmap

View heatmap of alert states for one or all Services, Areas, Environment or Regions, and optionally show Service Names. The heatmap organizes CIs by one or all Groups. Each rectangle in the heatmap represents a Group, which are grouped by Area. The rectangle size represents the number of CIs in the Service; a larger size is a larger value.

Use the available drop-down menus or right-click to filter data shown in the display. Use the check-boxes  to include or exclude labels in the heatmap. Move your mouse over a rectangle to see additional information. Drill-down and investigate by clicking a rectangle in the heatmap to view details in the last display that was viewed under either the **Service Summary Views** or **Key Metrics Views**. For example, if the last selected display was the **"Service Summary"** display under **"Service Summary Views"** and you clicked on a rectangle in the **Group / Region Heatmap**, the details would display in the **Service Summary** display. If the last selected display was the **"Service KM Table"** display under **"Key Metrics Views"**, then clicking a rectangle in the **Group / Region Heatmap** displays the details in the **Service KM Table**.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Note:** The "Up" Arrow () opens the most recently viewed display under **"All Management Areas"**. For example, if the last viewed display under **All Management Areas** was **Area Table**, then clicking opens the **"Area Table"** display.

**Filter By:**

The following filtering options are typically included:

**Owner:** Choose an Owner to see metrics for Areas associated with that Owner.

**Area:** Choose an Area to see metrics for Groups associated with that Area and Owner.

**Group:** Choose a Group to see metrics for Services associated with that Group, Area and Owner.

**Service:** Choose a Service to see metrics for Environments associated with that Service, Group, Area and Owner.

**Env:** Choose an Environment to see metrics for Environments associated with that Service, Group, Area and Owner.

**Metric:**

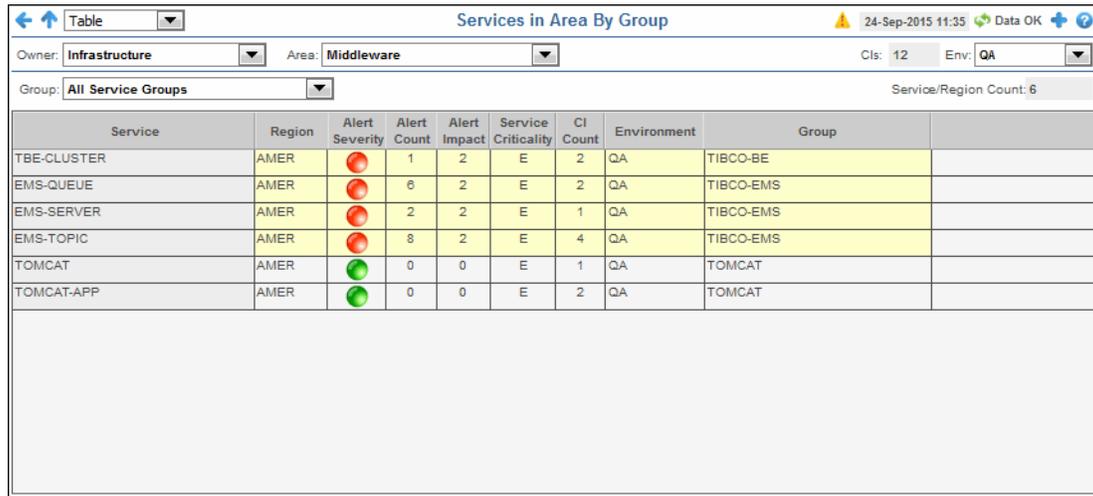
Choose the type of metric to show in the heatmap. Each metric has its own gradient bar that maps relative values to colors:

<b>Alert Impact</b>	The product of the maximum Alert Severity of alerts in the heatmap rectangle multiplied by the maximum Criticality of alerts in the heatmap rectangle. Values range from <b>0</b> - <b>10</b> , as indicated in the color gradient  bar, where <b>10</b> is the highest Alert Impact.
<b>Alert Severity</b>	The maximum level of alerts in the heatmap rectangle. Values range from <b>0</b> - <b>2</b> , as indicated in the color gradient  bar, where <b>2</b> is the highest Alert Severity. <ul style="list-style-type: none"> <li> Red indicates that one or more metrics have reached their alarm threshold. Metrics that have exceeded their specified ALARM LEVEL threshold have an Alert Severity value of <b>2</b>.</li> <li> Yellow indicates that one or more metrics have reached their alarm threshold. Metrics that have exceeded their specified WARNING LEVEL threshold have an Alert Severity value of <b>1</b>.</li> <li> Green indicates that no metrics have reached their alert thresholds. Metrics that have not exceeded their specified thresholds have an Alert Severity value of <b>0</b>.</li> </ul>
<b>Alert Count</b>	The total number of critical and warning alerts in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.
<b>Criticality</b>	The maximum level of Criticality (rank of importance) in the heatmap rectangle. Values range from <b>1</b> to <b>5</b> , as indicated in the color gradient  bar, where <b>5</b> is the highest Criticality. <p>Criticality is specified in the Service Data Model (CMDB) by your administrator. Criticality values are listed in the <b>Component Views</b> - "<a href="#">CI / Service Table</a>" display, which range from <b>A</b> to <b>E</b>, where <b>A</b> is the highest Criticality (level <b>5</b> maps to a Criticality of <b>A</b> and level <b>1</b> maps to a Criticality of <b>E</b> with equally spaced intermediate values).</p>

**Group / Service Table**

This table displays data shown in the **Group/Service** and **Group/Region** heatmaps. View Service metrics (Impact, Severity, Count and Criticality, and CI Count) for one or all Areas, Owners, Groups and Environments, and compare detailed metrics across all Areas in your organization. The table lists Services by Owner and Area. Each row in the table is a different Service. The color of the circle in the **Alert Severity** column represents the most critical alert state for that Service.

Use the available drop-down menus or right-click to filter data shown in the display. Click Sort  to order column data. Drill-down and investigate by clicking a row in the table to view details in the last display that was viewed under either the **Service Summary Views** or **Key Metric Views**. For example, if the last selected display was the “Service Summary” display under “Service Summary Views” and you clicked on a row in the table, the details would display in the **Service Summary** display. If the last selected display was the “Service KM Table” display under “Key Metrics Views”, then clicking a row in the table displays the details in the **Service KM Table**.



Service	Region	Alert Severity	Alert Count	Alert Impact	Service Criticality	CI Count	Environment	Group
TBE-CLUSTER	AMER		1	2	E	2	QA	TIBCO-BE
EMS-QUEUE	AMER		6	2	E	2	QA	TIBCO-EMS
EMS-SERVER	AMER		2	2	E	1	QA	TIBCO-EMS
EMS-TOPIC	AMER		8	2	E	4	QA	TIBCO-EMS
TOMCAT	AMER		0	0	E	1	QA	TOMCAT
TOMCAT-APP	AMER		0	0	E	2	QA	TOMCAT

#### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

The “Up” Arrow ( ) opens the most recently viewed display under “All Management Areas”. For example, if the last viewed display under **All Management Areas** was **Area Table**, then clicking opens the “Area Table” display.

#### Row Color Code:

Tables with colored rows indicate the following:

- Red indicates that one or more alerts exceeded their ALARM LEVEL threshold in the table row.
- Yellow indicates that one or more alerts exceeded their WARNING LEVEL threshold in the table row.
- Green indicates that no alerts exceeded their WARNING or ALARM LEVEL threshold in the table row.

**Filter By:**

The following filtering options are typically included:

**Owner:** Choose an Owner to see metrics for Areas associated with that Owner.

**Area:** Choose an Area to see metrics for Groups associated with that Area and Owner.

**Group:** Choose a Group to see metrics for Services associated with that Group, Area and Owner.

**Service:** Choose a Service to see metrics for Environments associated with that Service, Group, Area and Owner.

**Env:** Choose an Environment to see metrics for Environments associated with that Service, Group, Area and Owner.

**Fields and Data**

This display includes:

<b>Service/ Region Count</b>	The total number of Services listed in the table. This value is determined by the selections made from display drop-down menus.
<b>Area</b>	The name of the Area where the alert data originated.
<b>Service</b>	The name of the Service where the alert data originated.
<b>Region</b>	The name of the Region to which the Service applies.
<b>Severity</b>	The maximum level of alerts in the row. Values range from <b>0</b> to <b>2</b> , where <b>2</b> is the greatest Severity:  One or more alerts exceeded their ALARM LEVEL threshold in the Service.  One or more alerts exceeded their WARNING LEVEL threshold in the Service.  No alert thresholds have been exceeded in the Service.
<b>Alert Count</b>	The total number of critical and warning alerts for the Service.
<b>Alert Impact</b>	The maximum of the products of maximum Alert Severity multiplied by the Criticality of all CIs for the Service. Values range from <b>0</b> - <b>10</b> , where <b>10</b> is the highest Alert Impact.
<b>Service Criticality</b>	The Criticality (rank of importance) specified in the Service Data Model (CMDDB) by your administrator. Criticality values are listed in the <b>Component Views / CI Service Table</b> display, which range from A to E, where A is the highest Criticality.
<b>CIs</b>	The total number of configurable items in the display.
<b>Environment</b>	The name of the Environment to which the Service applies.
<b>Group</b>	The name of the Environment to which the Service applies.
<b>CI Count</b>	The total number of configurable items associated with the Service.

**Services CI Type Summary**

This display lists the health state of Services by CI Type and allows you to manage alerts. In the upper table, each column is a CI Type and each row is a Service. Select a row in the table to view details in the lower table.

Use the available drop-down menus or right-click to filter data shown in the display. Click Sort  to order column data.

**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu Table open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

**Note:** The “Up” Arrow () opens the most recently viewed display under “All Management Areas”. For example, if the last viewed display under **All Management Areas** was **Area Table**, then clicking  opens the “Area Table” display.

For each Service in a selected Group, the round indicator  shows the current maximum Alert Severity of all the CIs associated with each CI Type.

-  Red indicates that one or more alerts exceeded their ALARM LEVEL threshold in the table row.
-  Yellow indicates that one or more alerts exceeded their WARNING LEVEL threshold in the table row.
-  Green indicates that no alerts exceeded their WARNING or ALARM LEVEL threshold in the table row.

The cell background color indicates the current maximum Alert Impact of all the CIs associated with the Service and CI Type. The Alert Impact is calculated for each CI, which is the product of the CI Criticality times the current maximum Alert Severity. Background colors range from green to red, green being the lowest possible alert impact and red the highest possible value.

For example, in the following figure the first five Services in the list have an alert condition due to a BW Engine problem, and additionally the **INVENTORY MANAGER** Service has a TIBCO EMS Server problem. The **All CI Types** column shows the global highest level for all CI Types.

Service Name	All CI Types	User Experience	JVM	BW Server	BW Engine	TibcoEMS Server	TibcoEMS Topic	Tomcat
ACCOUNTING		0	0					
COMPLIANCE		0						
INVENTORY MANAGER		0	0					
ORDER PROCESSING		0	0					
REPORTING		0	0					
TUCON-EXCHANGE		0	0	0	0			0

**Filter By:**

The following filtering options are typically included:

**Owner:** Choose an Owner to see metrics for Areas associated with that Owner.

**Area:** Choose an Area to see metrics for Groups associated with that Area and Owner.

**Group:** Choose a Group to see metrics for Services associated with that Group, Area and Owner.

**Service:** Choose a Service to see metrics for Environments associated with that Service, Group, Area and Owner.

**Env:** Choose an Environment to see metrics for Environments associated with that Service, Group, Area and Owner.

**Fields and Data**

This display includes:

<b>Valid CI Types Only</b>	Check to only show CI Type columns that contain data in the table, uncheck to include columns that are empty. Including empty table columns can be helpful when you are comparing Services (using the <b>Group</b> drop-down menu) because the table columns retain their order.
<b>Service/Region Count</b>	The total number of Services currently listed in the table.
<b>Service Name</b>	The name of the Service.
<b>All CI Types</b>	The circular indicator  shows the current maximum Alert Severity of all the CIs associated with the CI Type, and the cell background color shows the current maximum Alert Impact of all the CIs--across all CI Types-- associated with the Service.
<b>Service</b>	Shows the Service selected in the upper table.
<b>CI Type</b>	Shows the CI Type selected in the upper table.

**Alerts Table**

This table lists all open, unsuppressed alerts associated with the selection in the upper table. Each row in the table is a different active alert. Select one or more rows, right-click to open the **Alert** popup menu and choose an action to perform on the alert(s): **Details, Own, Suppress, Close, Annotate** or **Options**. Use the sort  button to order column data. The row color indicates the following:

-  Red indicates that one or more alerts exceeded their ALARM LEVEL threshold in the table row.
  -  Yellow indicates that one or more alerts exceeded their WARNING LEVEL threshold in the table row.
  -  Green indicates that no alerts exceeded their WARNING or ALARM LEVEL threshold in the table row.
  -  Gray indicates that the alert engine that is hosting the alert is not connected, not enabled or not initialized. When you select a gray row the **Own, Suppress, Unsuppress, Close, Annotate, Options** and **Details** buttons are disabled.
-  Opens the **Alerts Table** display in a new window.

<b>Own</b>	Click to assign an Owner for the selected alert(s). This button is only visible to users with Administrator privileges. This button is disabled when you select a gray row.
<b>Suppress</b>	Click to suppress the selected alert(s). This button is only visible to users with Administrator privileges. This button is disabled when you select a gray row.
<b>Close</b>	Click to close the selected alert(s). This button is only visible to users with Administrator privileges. This button is disabled when you select a gray row.
<b>Details</b>	Select an alert, right-click and choose <b>Alert/Details</b> to open the <b>Alert Detail</b> window and view alert details. Or, double-click an alert to open the <b>Alert Detail</b> window.

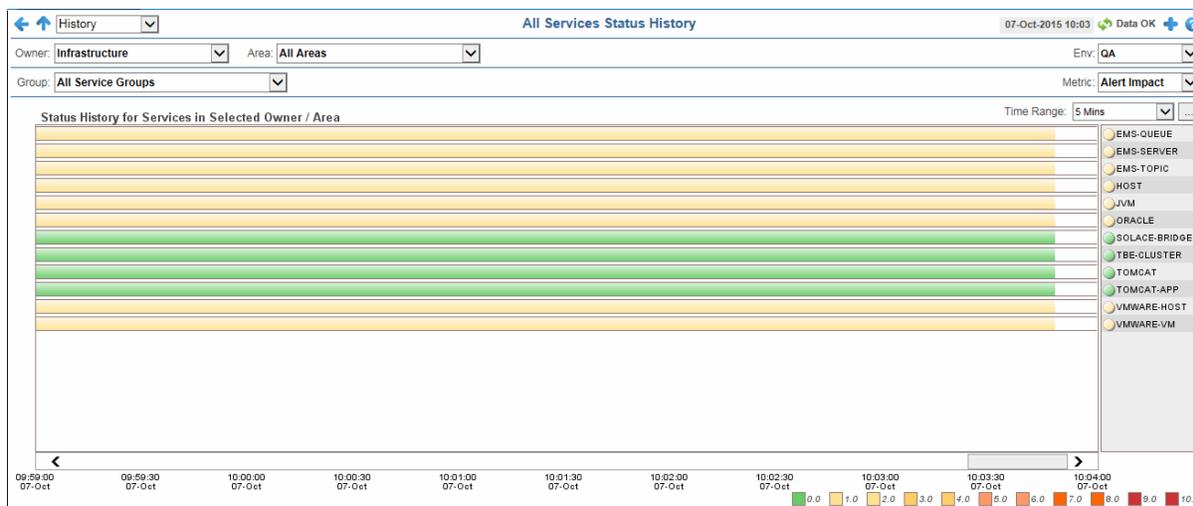
<b>Annotate</b>	Select one or more alerts, right-click and choose <b>Alert/Annotate</b> to open the <b>Set Owner and Comments</b> dialog and enter comments or change alert owner.
<b>Options</b>	Select an alert, right-click and choose <b>Alert/Options</b> to open the <b>Alert Options</b> dialog. This dialog is provided for customizing your own alert options.
<b>First Occ</b>	The date and time the alert first occurred.
<b>Last Occ</b>	The date and time the alert last occurred.
<b>Count</b>	The number of times the alert was generated.
<b>Sup</b>	When checked, the alert has been suppressed by a user.
<b>Owner</b>	The named owner assigned by the administrator.
<b>Alert Name</b>	The name of the alert.
<b>Primary Service</b>	The name of the Service with which the alert is associated.
<b>CI</b>	The CI alert source.
<b>Alert Text</b>	Description of the alert.
<b>AlertClass</b>	An optional alert field which can be used when integrating with other alerting systems.
<b>CompID</b>	An optional alert field which can be used when integrating with other alerting systems.
<b>TicketID</b>	An optional alert field which can be used when integrating with other alerting systems.
<b>TicketGroup</b>	An optional alert field which can be used when integrating with other alerting systems.

## Services History Heatmap

View history heatmap of alert states, over time, for Services in one Area, filtered by Group and Environment.

The history heatmap displays Services from one or more Groups and Environments of a given Owner and Area. Each row in the heatmap represents a different Service. The row color shows the Alert Impact or Alert Severity of a Service across time.

Use the available drop-down menus or right-click to filter data shown in the display. Mouse-over each row to see the time of alert state changes for particular Service occurred. For example, you can see at what time an alert state changed from green to red. Use the checkboxes  to include or exclude labels in the heatmap. Drill-down and investigate by clicking a row in the heatmap to view details in the last display that was viewed under either the **Service Summary Views** or **Key Metric Views**. For example, if the last selected display was the "Service Summary" display under "Service Summary Views" and you clicked on a row in the **Services History Heatmap**, the details would display in the **Service Summary** display. If the last selected display was the "Service KM Table" display under "Key Metrics Views", then clicking a row in the **Services History Heatmap** displays the details in the **Service KM Table**.



#### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu , Table open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

**Note:** The "Up" Arrow () opens the most recently viewed display under "All Management Areas". For example, if the last viewed display under **All Management Areas** was **Area Table**, then clicking opens the "Area Table" display.

#### Filter By:

The following filtering options are typically included:

**Owner:** Choose an Owner to see metrics for Areas associated with that Owner.

**Area:** Choose an Area to see metrics for Groups associated with that Area and Owner.

**Group:** Choose a Group to see metrics for Services associated with that Group, Area and Owner.

**Service:** Choose a Service to see metrics for Environments associated with that Service, Group, Area and Owner.

**Env:** Choose an Environment to see metrics for Environments associated with that Service, Group, Area and Owner.

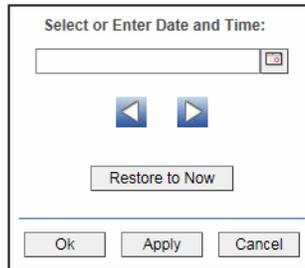
**Color Code:**

Row color indicates the following:

- Red indicates that one or more alerts exceeded their ALARM LEVEL threshold in the row.
- Yellow indicates that one or more alerts exceeded their WARNING LEVEL threshold in the row.
- Green indicates that no alerts exceeded their WARNING or ALARM LEVEL threshold in the row.

**Time Range**

Select a time range from the drop down menu varying from 2 Minutes to Last 7 Days, or display All Data. By default, the time range end point is the current time.



To change the time range for the graph, click Open Calendar , choose the date and time, then click **OK**. Or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the Time Range drop-down menu. Click Restore to Now to reset the time range end point to the current time.

## Single Area Service Views

These displays present aggregated data of alert states for all Services for a specific Area. Data can be filtered by Area, Group, Environment, and alert Metric. Data is filtered by the \$rtvOwnerMask, \$rtvAreaMask and \$rtvGroupMask values for the logged in user. For details, refer to the *RTView Enterprise Configuration Guide*.

Use these displays, for example, to isolate the Area and Environment in which a critical alert is occurring. If you see a critical alert, get information by comparing alert metrics (such as how many other items are potentially affected).

These displays drill-down to the **Service Summary Views - "Service By CI Type"** display. The **Single Area Service Views** displays present data in tabular and heatmap formats. Displays in this View are:

- "Single Area: Group/Service Heatmap": Heatmap of alert states for Services by Area, with the option to filter by Area, Group, Environment and alert Metric, and the option to show Group and Service Names.
- "Single Area: Region/Service Heatmap": Heatmap as described for the **Group / Service Heatmap** (above), with the option to filter by Region and no option to show Service Names.
- "Single Area: Group / Service Table": Table view of **Group/Service Heatmap** data.
- "Single Area: Services CI Type Summary": Table that shows the health state of Services per CI Type.
- "Single Area: Services History Heatmap": Heatmap of alert states, over time, for Services in a selected Area, with the option to filter by Group, Environment and alert Metric.

## Single Area: Group/Service Heatmap

View heatmap of alert states for Services in one Area, filter by Group or Environment, and optionally show Service Names. Each rectangle in the heatmap represents a Service (for example, Applications, Demo Systems and so forth), which are grouped by Area. The rectangle size represents the number of CIs in the Service; a larger size is a larger value.

Use the available drop-down menus or right-click to filter data shown in the display. Use the check-boxes  to include or exclude labels in the heatmap. Move your mouse over a rectangle to see additional information. Drill-down and investigate by clicking a rectangle in the heatmap to view details in the last display that was viewed under either the **Service Summary Views** or **Key Metrics Views**. For example, if the last selected display was the "Service Summary" display under "Service Summary Views" and you clicked on a rectangle in the **Group / Service Heatmap**, the details would display in the **Service Summary** display. If the last selected display was the "Service KM Table" display under "Key Metrics Views", then clicking a rectangle in the **Group / Service Heatmap** displays the details in the **Service KM Table**.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Note:** The "Up" Arrow () opens the most recently viewed display under "Multi Area Service Views". For example, if the last viewed display under **Multi Area Service Views** was **Group/Region Heatmap**, then clicking opens the "Group/Region Heatmap" display.

**Filter By:**

The following filtering options are typically included:

**Owner:** Choose an Owner to see metrics for Areas associated with that Owner.

**Area:** Choose an Area to see metrics for Groups associated with that Area and Owner.

**Group:** Choose a Group to see metrics for Services associated with that Group, Area and Owner.

**Service:** Choose a Service to see metrics for Environments associated with that Service, Group, Area and Owner.

**Env:** Choose an Environment to see metrics for Environments associated with that Service, Group, Area and Owner.

**Metric:**

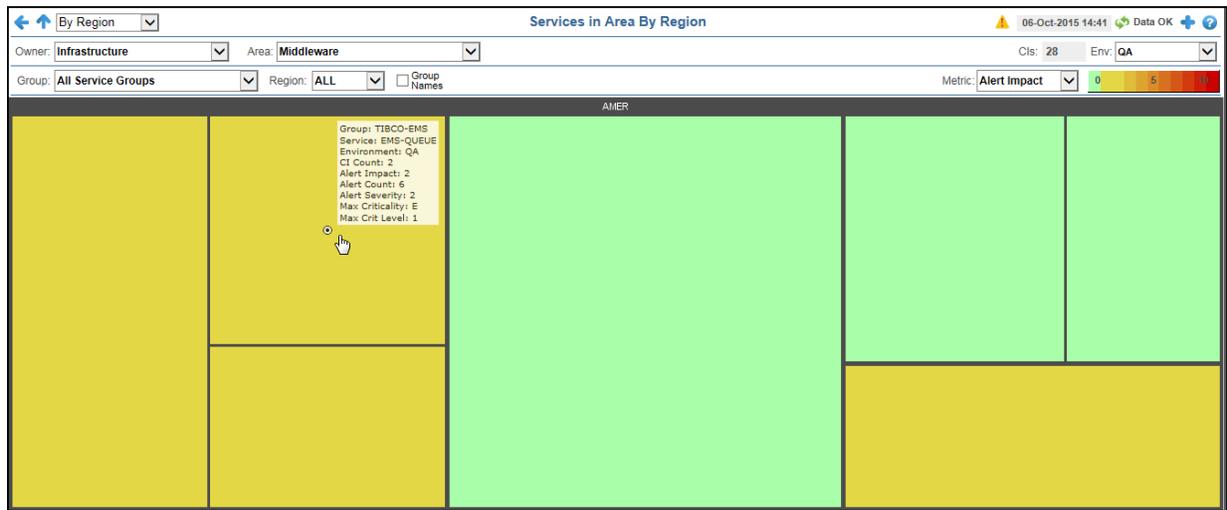
Choose the type of metric to show in the heatmap. Each metric has its own gradient bar that maps relative values to colors:

<b>Alert Impact</b>	The product of the maximum Alert Severity of alerts in the heatmap rectangle multiplied by the maximum Criticality of alerts in the heatmap rectangle. Values range from <b>0</b> - <b>10</b> , as indicated in the color gradient  bar, where <b>10</b> is the highest Alert Impact.
<b>Alert Severity</b>	The maximum level of alerts in the heatmap rectangle. Values range from <b>0</b> - <b>2</b> , as indicated in the color gradient  bar, where <b>2</b> is the highest Alert Severity. <ul style="list-style-type: none"> <li><span style="color: red;">●</span> Red indicates that one or more metrics have reached their alarm threshold. Metrics that have exceeded their specified ALARM LEVEL threshold have an Alert Severity value of <b>2</b>.</li> <li><span style="color: yellow;">●</span> Yellow indicates that one or more metrics have reached their alarm threshold. Metrics that have exceeded their specified WARNING LEVEL threshold have an Alert Severity value of <b>1</b>.</li> <li><span style="color: green;">●</span> Green indicates that no metrics have reached their alert thresholds. Metrics that have not exceeded their specified thresholds have an Alert Severity value of <b>0</b>.</li> </ul>
<b>Alert Count</b>	The total number of critical and warning alerts in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.
<b>Criticality</b>	The maximum level of Criticality (rank of importance) in the heatmap rectangle. Values range from <b>1</b> to <b>5</b> , as indicated in the color gradient  bar, where <b>5</b> is the highest Criticality. <p>Criticality is specified in the Service Data Model (CMDB) by your administrator. Criticality values are listed in the <b>Component Views</b> - "<a href="#">CI / Service Table</a>" display, which range from <b>A</b> to <b>E</b>, where <b>A</b> is the highest Criticality (level <b>5</b> maps to a Criticality of <b>A</b> and level <b>1</b> maps to a Criticality of <b>E</b> with equally spaced intermediate values).</p>

**Single Area: Region/Service Heatmap**

View heatmap of alert states for one Owner, one specific Area, one or all Service Groups, and one or all Regions. You can also optionally show Service Group Names. The heatmap organizes CIs by one or all Groups. Each rectangle in the heatmap represents a Group, which is grouped by Area. The rectangle size represents the number of CIs in the Service; a larger size is a larger value.

Use the available drop-down menus or right-click to filter data shown in the display. Use the check-boxes  to include or exclude labels in the heatmap. Move your mouse over a rectangle to see additional information. Drill-down and investigate by clicking a rectangle in the heatmap to view details in the last display that was viewed under either the **Service Summary Views** or **Key Metrics Views**. For example, if the last selected display was the “Service Summary” display under “Service Summary Views” and you clicked on a rectangle in the **Group / Region Heatmap**, the details would display in the **Service Summary** display. If the last selected display was the “Service KM Table” display under “Key Metrics Views”, then clicking a rectangle in the **Group / Region Heatmap** displays the details in the **Service KM Table**.



#### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** , **Table** open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Note:** The “Up” Arrow () opens the most recently viewed display under “Multi Area Service Views”. For example, if the last viewed display under Multi Area Service Views was Group/Region Heatmap, then clicking opens the “Group/Region Heatmap” display.

#### Filter By:

The following filtering options are typically included:

**Owner:** Choose an Owner to see metrics for Areas associated with that Owner.

**Area:** Choose an Area to see metrics for Groups associated with that Area and Owner.

**Group:** Choose a Group to see metrics for Services associated with that Group, Area and Owner.

**Service:** Choose a Service to see metrics for Environments associated with that Service, Group, Area and Owner.

**Env:** Choose an Environment to see metrics for Environments associated with that Service, Group, Area and Owner.

**Metric:**

Choose the type of metric to show in the heatmap. Each metric has its own gradient bar that maps relative values to colors:

<b>Alert Impact</b>	The product of the maximum Alert Severity of alerts in the heatmap rectangle multiplied by the maximum Criticality of alerts in the heatmap rectangle. Values range from <b>0</b> - <b>10</b> , as indicated in the color gradient  bar, where <b>10</b> is the highest Alert Impact.
<b>Alert Severity</b>	<p>The maximum level of alerts in the heatmap rectangle. Values range from <b>0</b> - <b>2</b>, as indicated in the color gradient  bar, where <b>2</b> is the highest Alert Severity.</p> <ul style="list-style-type: none"> <li> Red indicates that one or more metrics have reached their alarm threshold. Metrics that have exceeded their specified ALARM LEVEL threshold have an Alert Severity value of <b>2</b>.</li> <li> Yellow indicates that one or more metrics have reached their alarm threshold. Metrics that have exceeded their specified WARNING LEVEL threshold have an Alert Severity value of <b>1</b>.</li> <li> Green indicates that no metrics have reached their alert thresholds. Metrics that have not exceeded their specified thresholds have an Alert Severity value of <b>0</b>.</li> </ul>
<b>Alert Count</b>	The total number of critical and warning alerts in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.
<b>Criticality</b>	<p>The maximum level of Criticality (rank of importance) in the heatmap rectangle. Values range from <b>1</b> to <b>5</b>, as indicated in the color gradient  bar, where <b>5</b> is the highest Criticality.</p> <p>Criticality is specified in the Service Data Model (CMDB) by your administrator. Criticality values are listed in the <b>Component Views</b> - "<a href="#">CI / Service Table</a>" display, which range from <b>A</b> to <b>E</b>, where <b>A</b> is the highest Criticality (level <b>5</b> maps to a Criticality of <b>A</b> and level <b>1</b> maps to a Criticality of <b>E</b> with equally spaced intermediate values).</p>

## Single Area: Group / Service Table

This table displays data shown in the **Group/Service** and **Region/Service** heatmaps. View Service metrics (Impact, Severity, Count and Criticality, and CI Count) for a specific Area, for one or all Owners, Groups, and Environments. The table lists Services by Owner and Area. Each row in the table is a different Service. The color of the circle in the **Alert Severity** column represents the most critical alert state for that Service.

Use the available drop-down menus or right-click to filter data shown in the display. Click Sort  to order column data. Drill-down and investigate by clicking a row in the table to view details in the last display that was viewed under either the **Service Summary Views** or **Key Metric Views**. For example, if the last selected display was the “Service Summary” display under “Service Summary Views” and you clicked on a row in the table, the details would display in the **Service Summary** display. If the last selected display was the “Service KM Table” display under “Key Metrics Views”, then clicking a row in the table displays the details in the **Service KM Table**.

Service	Region	Alert Severity	Alert Count	Alert Impact	Service Criticality	CI Count	Environment	Group
SOLACE-BRIDGE	AMER		0	0	E	16	QA	SOLACE
TBE-CLUSTER	AMER		1	1	E	2	QA	TIBCO-BE
EMS-QUEUE	AMER		6	2	E	2	QA	TIBCO-EMS
EMS-SERVER	AMER		2	2	E	1	QA	TIBCO-EMS
EMS-TOPIC	AMER		8	2	E	4	QA	TIBCO-EMS
TOMCAT	AMER		0	0	E	1	QA	TOMCAT
TOMCAT-APP	AMER		0	0	E	2	QA	TOMCAT

#### Title Bar (possible features are):

-   Open the previous and upper display.
-  Open an instance of this display in a new window.
-  Open the online help page for this display.
- open commonly accessed displays.

 **Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Note:** The “Up” Arrow () opens the most recently viewed display under “Multi Area Service Views”. For example, if the last viewed display under **Multi Area Service Views** was **Group/Region Heatmap**, then clicking  opens the “Group/Region Heatmap” display.

#### Row Color Code:

Tables with colored rows indicate the following:

-  Red indicates that one or more alerts exceeded their ALARM LEVEL threshold in the table row.
-  Yellow indicates that one or more alerts exceeded their WARNING LEVEL threshold in the table row.
-  Green indicates that no alerts exceeded their WARNING or ALARM LEVEL threshold in the table row.

**Filter By:**

The following filtering options are typically included:

**Owner:** Choose an Owner to see metrics for Areas associated with that Owner.

**Area:** Choose an Area to see metrics for Groups associated with that Area and Owner.

**Group:** Choose a Group to see metrics for Services associated with that Group, Area and Owner.

**Service:** Choose a Service to see metrics for Environments associated with that Service, Group, Area and Owner.

**Env:** Choose an Environment to see metrics for Environments associated with that Service, Group, Area and Owner.

**Fields and Data**

This display includes:

<b>Service/ Region Count</b>	The total number of Services listed in the table. This value is determined by the selections made from display drop-down menus.
<b>Area</b>	The name of the Area where the alert data originated.
<b>Service</b>	The name of the Service where the alert data originated.
<b>Region</b>	The name of the Region to which the Service applies.
<b>Severity</b>	The maximum level of alerts in the row. Values range from <b>0</b> to <b>2</b> , where <b>2</b> is the greatest Severity: <ul style="list-style-type: none"> <li><span style="color: red;">●</span> One or more alerts exceeded their ALARM LEVEL threshold in the Service.</li> <li><span style="color: yellow;">●</span> One or more alerts exceeded their WARNING LEVEL threshold in the Service.</li> <li><span style="color: green;">●</span> No alert thresholds have been exceeded in the Service.</li> </ul>
<b>Alert Count</b>	The total number of critical and warning alerts for the Service.
<b>Alert Impact</b>	The maximum of the products of maximum Alert Severity multiplied by the Criticality of all CIs for the Service. Values range from <b>0</b> - <b>10</b> , where <b>10</b> is the highest Alert Impact.
<b>Service Criticality</b>	The Criticality (rank of importance) specified in the Service Data Model (CMDB) by your administrator. Criticality values are listed in the <b>Component Views / CI Service Table</b> display, which range from A to E, where A is the highest Criticality.
<b>CIs</b>	The total number of configurable items in the display.
<b>Environment</b>	The name of the Environment to which the Service applies.
<b>Group</b>	The name of the Environment to which the Service applies.
<b>CI Count</b>	The total number of configurable items associated with the Area.

**Single Area: Services CI Type Summary**

This display lists the health state of Services by CI Type and allows you to manage alerts. In the upper table, each column is a CI Type and each row is a Service. Select a row in the table to view details in the lower table.

Use the available drop-down menus or right-click to filter data shown in the display. Click Sort  to order column data.

Service Health By CI Type

Owner: Infrastructure Area: Middleware Cts: 28 Env: QA

Group: All Service Groups Valid CI Types Only Service/Region Count: 7

Service Name	All CI Types	AMX Node	AMX Service	AMX ServiceNod	BW Server	BW Engine	BW Proc	EMS Server	EMS Topic	EMS Queue	Active Spaces	Tomcat	Tomcat App	ServiceGroup
SOLACE-BRIDGE		0	0	0	0	0	0	0	0	0	0	0	0	SOLACE
TBE-CLUSTER		0	0	0	0	0	0	0	0	0	0	0	0	TIBCO-BE
EMS-QUEUE		0	0	0	0	0	0	0	0		0	0	0	TIBCO-EMS
EMS-SERVER		0	0	0	0	0	0		0	0	0	0	0	TIBCO-EMS
EMS-TOPIC		0	0	0	0	0	0	0		0	0	0	0	TIBCO-EMS
TOMCAT		0	0	0	0	0	0	0	0	0	0		0	TOMCAT
TOMCAT-APP		0	0	0	0	0	0	0	0	0	0		0	TOMCAT

Service: EMS-QUEUE CI Type: \*

First Occ	Last Occ	Count	Sup	Owner	Alert Name	Primary Service	CI	Alert Text
10/06/15 06:18:02	10/06/15 06:18:02	1	<input type="checkbox"/>		EmsQueueProviderIdle...	EMS	tcp://192.168.200.13...	High Alert Limit exceeded, current value: 97.0 limit: 80.0
10/06/15 06:18:02	10/06/15 06:18:02	1	<input type="checkbox"/>		EmsQueueProviderIdle...	EMS	tcp://192.168.200.13...	High Alert Limit exceeded, current value: 97.0 limit: 80.0
10/06/15 06:16:48	10/06/15 06:16:48	1	<input type="checkbox"/>		EmsQueuesProducerC...	EMS	tcp://192.168.200.13...	Low Alert Limit exceeded, current value: 0.0 limit: 5.0
10/06/15 06:16:48	10/06/15 06:16:48	1	<input type="checkbox"/>		EmsQueuesProducerC...	EMS	tcp://192.168.200.13...	Low Alert Limit exceeded, current value: 0.0 limit: 5.0
10/06/15 06:16:48	10/06/15 06:16:48	1	<input type="checkbox"/>		EmsQueuesConsumer...	EMS	tcp://192.168.200.13...	Low Alert Limit exceeded, current value: 0.0 limit: 5.0
10/06/15 06:16:48	10/06/15 06:16:48	1	<input type="checkbox"/>		EmsQueuesConsumer...	EMS	tcp://192.168.200.13...	Low Alert Limit exceeded, current value: 0.0 limit: 5.0

**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu Table open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

**Note:** The "Up" Arrow ( ) opens the most recently viewed display under "Multi Area Service Views". For example, if the last viewed display under **Multi Area Service Views** was **Group/Region Heatmap**, then clicking opens the "Group/Region Heatmap" display.

For each Service in a selected Group, the round indicator shows the current maximum Alert Severity of all the CIs associated with each CI Type.

- Red indicates that one or more alerts exceeded their ALARM LEVEL threshold in the table row.
- Yellow indicates that one or more alerts exceeded their WARNING LEVEL threshold in the table row.
- Green indicates that no alerts exceeded their WARNING or ALARM LEVEL threshold in the table row.

The cell background color indicates the current maximum Alert Impact of all the CIs associated with the Service and CI Type. The Alert Impact is calculated for each CI, which is the product of the CI Criticality times the current maximum Alert Severity. Background colors range from green to red, green being the lowest possible alert impact and red the highest possible value.

For example, in the following figure the first five Services in the list have an alert condition due to a BW Engine problem, and additionally the **INVENTORY MANAGER** Service has a TIBCO EMS Server problem. The **All CI Types** column shows the global highest level for all CI Types.

Service Name	All CI Types	User Experience	JVM	BW Server	BW Engine	TibcoEMS Server	TibcoEMS Topic	Tomcat
ACCOUNTING		0	0					
COMPLIANCE		0	0					
INVENTORY MANAGER		0	0					
ORDER PROCESSING		0	0					
REPORTING			0					
TUCON-EXCHANGE		0	0	0	0			0

**Filter By:**

The following filtering options are typically included:

**Owner:** Choose an Owner to see metrics for Areas associated with that Owner.

**Area:** Choose an Area to see metrics for Groups associated with that Area and Owner.

**Group:** Choose a Group to see metrics for Services associated with that Group, Area and Owner.

**Service:** Choose a Service to see metrics for Environments associated with that Service, Group, Area and Owner.

**Env:** Choose an Environment to see metrics for Environments associated with that Service, Group, Area and Owner.

**Fields and Data**

This display includes:

- Valid CI Types Only** Check to only show CI Type columns that contain data in the table, uncheck to include columns that are empty. Including empty table columns can be helpful when you are comparing Services (using the **Group** drop-down menu) because the table columns retain their order.
- Service / Region Count** The total number of Services currently listed in the table.
- Service Name** The name of the Service.
- All CI Types** The circular indicator  shows the current maximum Alert Severity of all the CIs associated with the CI Type, and the cell background color shows the current maximum Alert Impact of all the CIs--across all CI Types-- associated with the Service.
- Service** Shows the Service selected in the upper table.
- CI Type** Shows the CI Type selected in the upper table.

**Alerts Table**

This table lists all open, unsuppressed alerts associated with the selection in the upper table. Each row in the table is a different active alert. Select one or more rows, right-click to open the **Alert** popup menu and choose an action to perform on the alert(s): **Details, Own, Suppress, Close, Annotate** or **Options**. Use the sort  button to order column data. The row color indicates the following:

-  Red indicates that one or more alerts exceeded their ALARM LEVEL threshold in the table row.
-  Yellow indicates that one or more alerts exceeded their WARNING LEVEL threshold in the table row.
-  Green indicates that no alerts exceeded their WARNING or ALARM LEVEL threshold in the table row.
-  Gray indicates that the alert engine that is hosting the alert is not connected, not enabled or not initialized. When you select a gray row the **Own, Suppress, Unsuppress, Close, Annotate, Options** and **Details** buttons are disabled.
-  Opens the **Alerts Table** display in a new window.

- Own** Click to assign an Owner for the selected alert(s). This button is only visible to users with Administrator privileges. This button is disabled when you select a gray row.
- Suppress** Click to suppress the selected alert(s). This button is only visible to users with Administrator privileges. This button is disabled when you select a gray row.
- Close** Click to close the selected alert(s). This button is only visible to users with Administrator privileges. This button is disabled when you select a gray row.

<b>Details</b>	Select an alert, right-click and choose <b>Alert/Details</b> to open the <b>Alert Detail</b> window and view alert details. Or, double-click an alert to open the <b>Alert Detail</b> window.
<b>Annotate</b>	Select one or more alerts, right-click and choose <b>Alert/Annotate</b> to open the <b>Set Owner and Comments</b> dialog and enter comments or change alert owner.
<b>Options</b>	Select an alert, right-click and choose <b>Alert/Options</b> to open the <b>Alert Options</b> dialog. This dialog is provided for customizing your own alert options.
<b>First Occ</b>	The date and time the alert first occurred.
<b>Last Occ</b>	The date and time the alert last occurred.
<b>Count</b>	The number of times the alert was generated.
<b>Sup</b>	When checked, the alert has been suppressed by a user.
<b>Owner</b>	The named owner assigned by the administrator.
<b>Alert Name</b>	The name of the alert.
<b>Primary Service</b>	The name of the Service with which the alert is associated.
<b>CI</b>	The CI alert source.
<b>Alert Text</b>	Description of the alert.
<b>AlertClass</b>	An optional alert field which can be used when integrating with other alerting systems.
<b>CompID</b>	An optional alert field which can be used when integrating with other alerting systems.
<b>TicketID</b>	An optional alert field which can be used when integrating with other alerting systems.
<b>TicketGroup</b>	An optional alert field which can be used when integrating with other alerting systems.

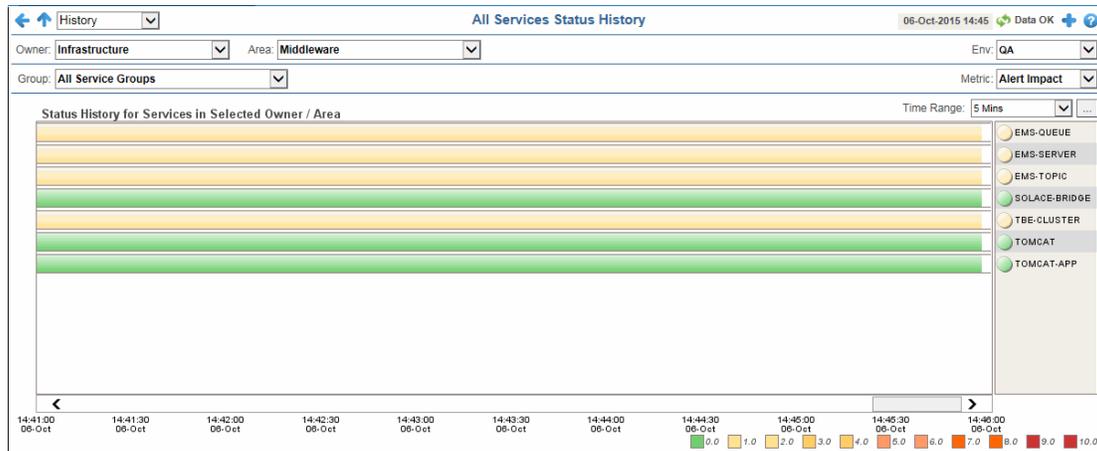
## Single Area: Services History Heatmap

View history heatmap of alert states, over time, for Services in one Area, filtered by Group and Environment.

The history heatmap displays Services from one or more Groups and Environments of a given Owner and Area. Each row in the heatmap represents a different Service. The row color shows the Alert Impact or Alert Severity of a Service across time.

Use the available drop-down menus or right-click to filter data shown in the display. Mouse-over each row to see the time of alert state changes for particular Service occurred. For example, you can see at what time an alert state changed from green to red. Use the checkboxes  to include or exclude labels in the heatmap. Drill-down and investigate by clicking a row in the heatmap to view details in the last display that was viewed under either the **Service Summary Views** or **Key Metric Views**.

For example, if the last selected display was the "Service KM Table" display under "Key Metrics Views" and you clicked on a row in the table, the details would display in the **Service Summary** display. If the last selected display was the "Service KM Table" display under "Key Metrics Views", then clicking a row in the table displays the details in the **Service KM Table**.



#### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Note:** The "Up" Arrow () opens the most recently viewed display under "Multi Area Service Views". For example, if the last viewed display under **Multi Area Service Views** was **Group/Region Heatmap**, then clicking opens the "Group/Region Heatmap" display.

#### Filter By:

The following filtering options are typically included:

- Owner:** Choose an Owner to see metrics for Areas associated with that Owner.
- Area:** Choose an Area to see metrics for Groups associated with that Area and Owner.
- Group:** Choose a Group to see metrics for Services associated with that Group, Area and Owner.
- Service:** Choose a Service to see metrics for Environments associated with that Service, Group, Area and Owner.
- Env:** Choose an Environment to see metrics for Environments associated with that Service, Group, Area and Owner.

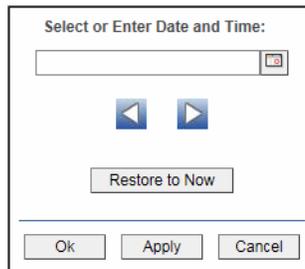
#### Color Code:

Row color indicates the following:

- Red indicates that one or more alerts exceeded their ALARM LEVEL threshold in the row.
- Yellow indicates that one or more alerts exceeded their WARNING LEVEL threshold in the row.
- Green indicates that no alerts exceeded their WARNING or ALARM LEVEL threshold in the row.

**Time Range**

Select a time range from the drop down menu varying from 2 Minutes to Last 7 Days, or display All Data. By default, the time range end point is the current time.



To change the time range for the graph, click Open Calendar , choose the date and time, then click **OK**. Or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the Time Range drop-down menu. Click Restore to Now to reset the time range end point to the current time.

## Service Summary Views

These displays present alert states at the component-level by Service in tabular and heatmap formats, while highlighting the most critical alert state. Data can be filtered by Owner, Area, Group, Service or Environment. Data is filtered by the \$rtvOwnerMask, \$rtvAreaMask, \$rtvGroupMask and \$rtvServiceMask values for the logged in user. For details, refer to the *RTView Enterprise Configuration Guide*.

Use these displays to get alert details and detailed status information for a particular Service, such as a list of all the CI Types relevant to a Service and the quality of the performance metrics for each CI. Displays in this View are:

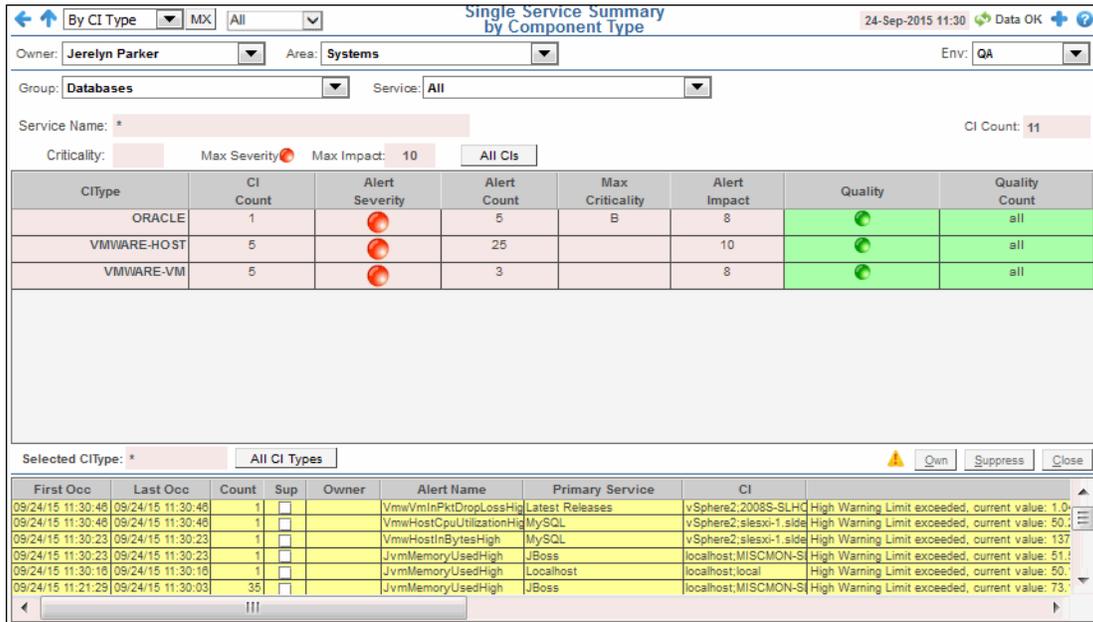
- "Service By CI Type": Table of alert states for a Service organized CI Type, with general alert information.
- "Service Summary": Table of CIs by Service, with detailed alert information.
- "Service Health Heatmap": Heatmap of CIs by Service, with the option to filter by Owner, Area, Group, Environment and alert Metric, and show CI Names.

### Service By CI Type

View alert states for a Service organized CI Type and manage alerts. See the CI Count for a Service and obtain alert statistics for CI Types such as Alert Severity and Alert Count. Use this display to summarize alerts occurring for a Service and determine which component types are malfunctioning. View a list of all active alerts associated with the CI Type.

The upper table lists all CI Types for the selected Service with alert details such as the highest Alert Severity. Each row is a CI Type. The color of each row represents the maximum Alert Impact for the row. Select a row that has an active alert (the Alert Severity is red or yellow) to view the active alerts in the lower table. Double-click a row to view a detailed list of CIs associated with the CI Type in the **Service Summary** display. In the lower table, each row is a different alert for a CI that is associated with the CI Type selected from the upper table.

Use the available drop-down menus or right-click to filter data shown in the display. Click Sort  to order column data.



CIType	CI Count	Alert Severity	Alert Count	Max Criticality	Alert Impact	Quality	Quality Count
ORACLE	1		5	B	8		all
VMWARE-HOST	5		25		10		all
VMWARE-VM	5		3		8		all

First Occ	Last Occ	Count	Sup	Owner	Alert Name	Primary Service	CI
09/24/15 11:30:48	09/24/15 11:30:48	1	<input type="checkbox"/>		VmwVmInPktDropLossHigh	Latest Releases	vSphere2:2008S-SLHC
09/24/15 11:30:48	09/24/15 11:30:48	1	<input type="checkbox"/>		VmwHostCpuUtilizationHigh	MySQL	vSphere2:slesxi-1.slide
09/24/15 11:30:23	09/24/15 11:30:23	1	<input type="checkbox"/>		VmwHostInBytesHigh	MySQL	vSphere2:slesxi-1.slide
09/24/15 11:30:23	09/24/15 11:30:23	1	<input type="checkbox"/>		JvmMemoryUsedHigh	JBoss	localhost:MISCMON-S
09/24/15 11:30:18	09/24/15 11:30:18	1	<input type="checkbox"/>		JvmMemoryUsedHigh	Localhost	localhost:local
09/24/15 11:21:29	09/24/15 11:30:03	35	<input type="checkbox"/>		JvmMemoryUsedHigh	JBoss	localhost:MISCMON-S

#### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

- Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

**Note:** The "Up" Arrow ( ) opens the most recently viewed display under "Multi Area Service Views". For example, if the last viewed display under **Multi Area Service Views** was **Services CI Type Summary**, then clicking opens the "Services CI Type Summary" display.

#### Filter By:

The following filtering options are typically included:

**Owner:** Choose an Owner to see metrics for Areas associated with that Owner.

**Area:** Choose an Area to see metrics for Groups associated with that Area and Owner.

**Group:** Choose a Group to see metrics for Services associated with that Group, Area and Owner.

**Service:** Choose a Service to see metrics for Environments associated with that Service, Group, Area and Owner.

**Env:** Choose an Environment to see metrics for Environments associated with that Service, Group, Area and Owner.

#### Fields and Data

This display includes:

**Service Name** The name of the selected Service.

**CI Count** The total number of configurable items in the display.

<b>Criticality</b>	The Criticality (rank of importance) specified in the Service Data Model (CMDB) by your administrator. Criticality values are listed in the Component Views / CI Service Table display, which range from <b>A</b> to <b>E</b> , where <b>A</b> is the highest Criticality. This value is used to determine the value for Alert Impact.
<b>Max Severity</b>	The highest Alert Severity value of any CI associated with the selected Service. Values range from <b>0</b> to <b>2</b> , where <b>2</b> is the greatest Severity: <ul style="list-style-type: none"> <li> One or more alerts exceeded their ALARM LEVEL threshold in the Service.</li> <li> One or more alerts exceeded their WARNING LEVEL threshold in the Service.</li> <li> No alert thresholds have been exceeded in the Service.</li> </ul>
<b>Max Impact</b>	The highest Alert Impact value of any CI associated with the selected Service.
<b>All CIs</b>	Opens the <b>Service Summary</b> display.

**(CI Type Table)**

This table lists all CI Types for the selected Service. Each row in the table is a CI Type. Click a row to view details in the lower table about alerts associated with the CI Type. Double-click a row to drill-down to Service Summary display describing alert details relevant to this CI Type.

<b>CIType</b>	The type of CI.
<b>CI Count</b>	The total number of configurable items associated with the CI Type.
<b>Alert Severity</b>	The highest Alert Severity value of any CI associated with the selected Service. Values range from <b>0</b> to <b>2</b> , where <b>2</b> is the greatest Severity: <ul style="list-style-type: none"> <li> One or more alerts exceeded their ALARM LEVEL threshold.</li> <li> One or more alerts exceeded their WARNING LEVEL threshold.</li> <li> No alert thresholds have been exceeded.</li> </ul>
<b>Alert Count</b>	The total number of active alerts for the CIs associated with the CI Type.
<b>Quality</b>	Shows whether performance metrics are being received from the CIs associated with the CI Type. <ul style="list-style-type: none"> <li><input type="radio"/> One or more performance metrics are not being received from the CIs associated with the CI Type.</li> <li><input checked="" type="radio"/> All performance metrics are being received from the CIs associated with the CI Type.</li> </ul>
<b>Quality Count</b>	Shows the number of CIs for that CI Type that have a known state. It displays all when that number is the total count of CI's. <ul style="list-style-type: none"> <li><input type="radio"/> One or more performance metrics are not being received from the CIs associated with the CI Type.</li> <li><input checked="" type="radio"/> All performance metrics are being received from the CIs associated with the CI Type.</li> </ul>
<b>Selected CI Type</b>	Shows the CI Type selected in the upper table.
<b>All CI Types</b>	Shows all active alerts for all CIs associated with the CI Type selected.

### Alerts Table

This table lists all open, unsuppressed alerts associated with the selection in the upper table. Each row in the table is a different active alert. Select one or more rows, right-click to open the **Alert** popup menu and choose an action to perform on the alert(s): **Details, Own, Suppress, Close, Annotate** or **Options**. Use the sort  button to order column data. The row color indicates the following:

-  Red indicates that one or more alerts exceeded their ALARM LEVEL threshold in the table row.
-  Yellow indicates that one or more alerts exceeded their WARNING LEVEL threshold in the table row.
-  Green indicates that no alerts exceeded their WARNING or ALARM LEVEL threshold in the table row.
-  Gray indicates that the alert engine that is hosting the alert is not connected, not enabled or not initialized. When you select a gray row the **Own, Suppress, Unsuppress, Close, Annotate, Options** and **Details** buttons are disabled.
-  Opens the **Alerts Table** display in a new window.

- Own** Click to assign an Owner for the selected alert(s). This button is only visible to users with Administrator privileges. This button is disabled when you select a gray row.
- Suppress** Click to suppress the selected alert(s). This button is only visible to users with Administrator privileges. This button is disabled when you select a gray row.
- Close** Click to close the selected alert(s). This button is only visible to users with Administrator privileges. This button is disabled when you select a gray row.
- Details** Select an alert, right-click and choose **Alert/Details** to open the **Alert Detail** window and view alert details. Or, double-click an alert to open the **Alert Detail** window.
- Annotate** Select one or more alerts, right-click and choose **Alert/Annotate** to open the **Set Owner and Comments** dialog and enter comments or change alert owner.
- Options** Select an alert, right-click and choose **Alert/Options** to open the **Alert Options** dialog. This dialog is provided for customizing your own alert options.

<b>First Occ</b>	The date and time the alert first occurred.
<b>Last Occ</b>	The date and time the alert last occurred.
<b>Count</b>	The number of times the alert was generated.
<b>Sup</b>	When checked, the alert has been suppressed by a user.
<b>Owner</b>	The named owner assigned by the administrator.
<b>Alert Name</b>	The name of the alert.
<b>Primary Service</b>	The name of the Service with which the alert is associated.
<b>CI</b>	The CI alert source.
<b>Alert Text</b>	Description of the alert.
<b>AlertClass</b>	An optional alert field which can be used when integrating with other alerting systems.
<b>CompID</b>	An optional alert field which can be used when integrating with other alerting systems.
<b>TicketID</b>	An optional alert field which can be used when integrating with other alerting systems.
<b>TicketGroup</b>	An optional alert field which can be used when integrating with other alerting systems.

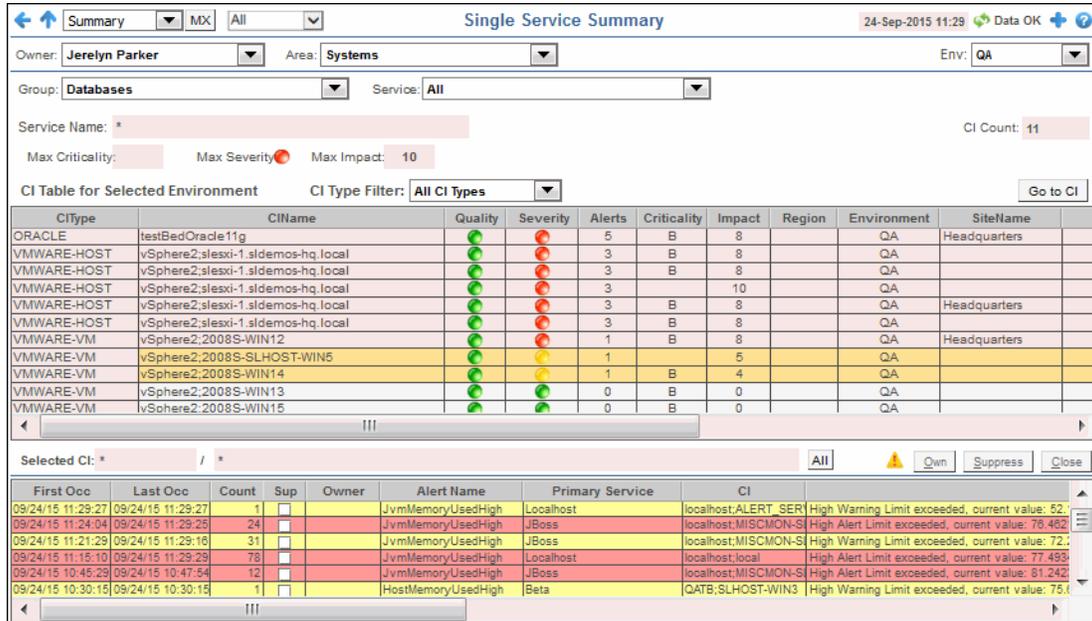
## Service Summary

View alert states at the component-level per Service, manage alerts, obtain component details such as the number of active alerts for the component, which operating system the component uses and the Data Server associated with the component.

Use this display to monitor a Service in a specific Group or Environment anywhere in your organization, and determine whether a component is malfunctioning.

The table lists all components for a selected Service. Each row in the table is a different CI (configurable item or component). Each CI can have multiple alerts. Click a row to view details in the lower table about any alerts associated with the CI.

Use the available drop-down menus or right-click to filter data shown in the display. Click Sort  to order column data. Double-click a row to drill-down to a summary page describing information relevant to this CI. This action can also be performed by selecting (a single click) on a row and selecting the **Go to CI** button.



The screenshot shows the 'Single Service Summary' interface. At the top, there are navigation arrows, a 'Summary' dropdown, and a 'Data OK' status indicator. Below this, the user's name 'Jeryln Parker' and the area 'Systems' are displayed. The 'Group' is set to 'Databases' and 'Service' to 'All'. A search bar for 'Service Name' is present, along with 'Max Criticality', 'Max Severity', and 'Max Impact' filters. A 'CI Table for Selected Environment' is shown with a 'CI Type Filter' set to 'All CI Types' and a 'Go to CI' button. The table lists various CI types such as ORACLE, VMWARE-HOST, and VMWARE-VM, with columns for Quality, Severity, Alerts, Criticality, Impact, Region, Environment, and SiteName. Below the table, there is a 'Selected CI' section with a search bar and buttons for 'All', 'Own', 'Suppress', and 'Close'. At the bottom, a detailed table shows alert history with columns for 'First Occ', 'Last Occ', 'Count', 'Sup', 'Owner', 'Alert Name', 'Primary Service', and 'CI'.

**Title Bar (possible features are):**

-  Open the previous and upper display.
-  Open an instance of this display in a new window.
-  Open the online help page for this display.
-   open commonly accessed displays.

 **Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

 **23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Note:** The "Up" Arrow () opens the most recently viewed display under "Multi Area Service Views". For example, if the last viewed display under **Multi Area Service Views** was **Services CI Type Summary**, then clicking  opens the "Services CI Type Summary" display.

**Row Color Code:**

Tables with colored rows indicate the following:

-  Red indicates that one or more alerts exceeded their ALARM LEVEL threshold in the table row.
-  Yellow indicates that one or more alerts exceeded their WARNING LEVEL threshold in the table row.
-  Green indicates that no alerts exceeded their WARNING or ALARM LEVEL threshold in the table row.

**Filter By:**

Use the available drop-down menus to filter data shown in the table. The display might include these filtering options:

**Owner:** Choose an Owner to see metrics in the heatmap for Areas associated with that Owner.

**Area:** Choose an Area to see metrics in the heatmap for Groups associated with that Area and Owner.

**Group:** Choose a Group to see metrics in the heatmap for Services associated with that Group, Area and Owner.

**Service:** Choose a Service to see metrics in the heatmap for Environments associated with that Service, Group, Area and Owner.

**Env:** Choose an Environment to see metrics in the heatmap for Environments associated with that Service, Group, Area and Owner.

**Fields and Data**

This display includes:

<b>Service Name</b>	The name of the selected Service.
<b>Criticality</b>	The Criticality (rank of importance) specified in the Service Data Model (CMDDB) by your administrator. Criticality values are listed in the Component Views / CI Service Table display, which range from <b>A</b> to <b>E</b> , where <b>A</b> is the highest Criticality. This value is used to determine the value for Alert Impact.
<b>CI Count</b>	The total number of configurable items in the display.

**CI Table for Selected Environment**

This table lists all CIs for the selected Environment. Each row in the table is a CI. Each CI can have multiple alerts. Click a row to view details about any alerts associated with the CI in the lower table. Double-click a row to drill-down to a summary page describing information relevant to this CI. This action can also be performed by selecting (a single click) on a row and selecting the **Go to CI** button.

<b>CI Type Filter</b>	Select a CI Type to display in the table or select All CI Types.
<b>Go to CI</b>	Drill-down to a summary page describing information relevant to this CI.
<b>CIType</b>	The type of CI.
<b>Quality</b>	Shows whether performance metrics are being received from the CI:  Performance metrics are not being received from the CI.  Performance metrics are being received from the CI.
<b>Severity</b>	Shows the most critical alert state for the selected CI:  One or more alerts exceeded their ALARM LEVEL threshold.  One or more alerts exceeded their WARNING LEVEL threshold.  No alert thresholds have been exceeded.
<b>Alerts</b>	The number of currently active alerts for the selected CI.
<b>Region</b>	The name of the Region for the CI.
<b>SiteName</b>	The name of the Site for the CI.
<b>OSType</b>	The operating system currently running on the CI.
<b>City</b>	The name of the City for the CI.
<b>Country</b>	The name of the Country for the CI.
<b>Data Server</b>	The name of the Data Server with which the CI is associated.

**Selected CI** Shows the CI Type selected in the upper table.

**All** Shows all active alerts for all CIs associated with the CI Type selected.

### Alerts Table

This table lists all open, unsuppressed alerts associated with the selection in the upper table. Each row in the table is a different active alert. Select one or more rows, right-click to open the **Alert** popup menu and choose an action to perform on the alert(s): **Details, Own, Suppress, Close, Annotate** or **Options**. Use the sort  button to order column data. The row color indicates the following:

 Red indicates that one or more alerts exceeded their ALARM LEVEL threshold in the table row.

 Yellow indicates that one or more alerts exceeded their WARNING LEVEL threshold in the table row.

 Green indicates that no alerts exceeded their WARNING or ALARM LEVEL threshold in the table row.

 Gray indicates that the alert engine that is hosting the alert is not connected, not enabled or not initialized. When you select a gray row the **Own, Suppress, Unsuppress, Close, Annotate, Options** and **Details** buttons are disabled.

 Opens the **Alerts Table** display in a new window.

**Own** Click to assign an Owner for the selected alert(s). This button is only visible to users with Administrator privileges. This button is disabled when you select a gray row.

**Suppress** Click to suppress the selected alert(s). This button is only visible to users with Administrator privileges. This button is disabled when you select a gray row.

**Close** Click to close the selected alert(s). This button is only visible to users with Administrator privileges. This button is disabled when you select a gray row.

**Details** Select an alert, right-click and choose **Alert/Details** to open the **Alert Detail** window and view alert details. Or, double-click an alert to open the **Alert Detail** window.

**Annotate** Select one or more alerts, right-click and choose **Alert/Annotate** to open the **Set Owner and Comments** dialog and enter comments or change alert owner.

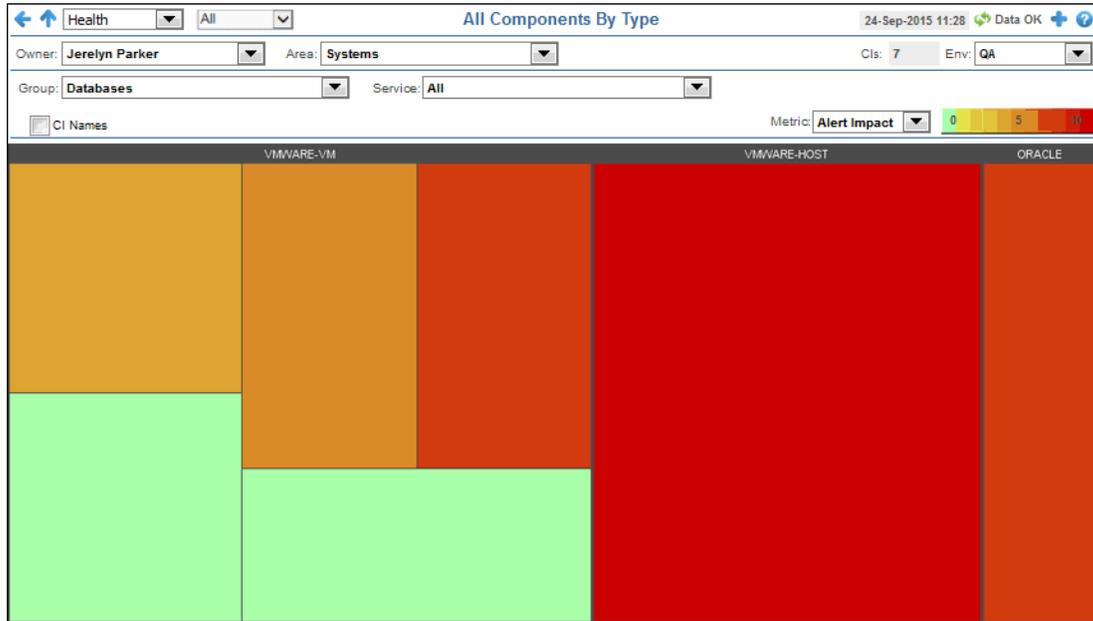
**Options** Select an alert, right-click and choose **Alert/Options** to open the **Alert Options** dialog. This dialog is provided for customizing your own alert options.

<b>First Occ</b>	The date and time the alert first occurred.
<b>Last Occ</b>	The date and time the alert last occurred.
<b>Count</b>	The number of times the alert was generated.
<b>Sup</b>	When checked, the alert has been suppressed by a user.
<b>Owner</b>	The named owner assigned by the administrator.
<b>Alert Name</b>	The name of the alert.
<b>Primary Service</b>	The name of the Service with which the alert is associated.
<b>CI</b>	The CI alert source.
<b>Alert Text</b>	Description of the alert.
<b>AlertClass</b>	An optional alert field which can be used when integrating with other alerting systems.
<b>CompID</b>	An optional alert field which can be used when integrating with other alerting systems.
<b>TicketID</b>	An optional alert field which can be used when integrating with other alerting systems.
<b>TicketGroup</b>	An optional alert field which can be used when integrating with other alerting systems.

## Service Health Heatmap

View heatmap of alert states for CIs associated with a Service. The heatmap organizes CIs by the Service selected. Each rectangle in the heatmap represents a CI (for example, **localhost-14**). Each Metric (selected from the drop-down menu) has a color gradient bar that maps relative values to colors.

Use the available drop-down menus or right-click to filter data shown in the display. Use the check-boxes  to include or exclude labels in the heatmap. Move your mouse over a rectangle to see additional information. By default, this display shows Alert Impact.



#### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

**Note:** The "Up" Arrow () opens the most recently viewed display under "Multi Area Service Views". For example, if the last viewed display under **Multi Area Service Views** was **Services CI Type Summary**, then clicking opens the "Services CI Type Summary" display.

#### Filter By:

The following filtering options are typically included:

**Owner:** Choose an Owner to see metrics for Areas associated with that Owner.

**Area:** Choose an Area to see metrics for Groups associated with that Area and Owner.

**Group:** Choose a Group to see metrics for Services associated with that Group, Area and Owner.

**Service:** Choose a Service to see metrics for Environments associated with that Service, Group, Area and Owner.

**Env:** Choose an Environment to see metrics for Environments associated with that Service, Group, Area and Owner.

**Metric:**

Choose the type of metric to show in the heatmap. Each metric has its own gradient bar that maps relative values to colors:

<b>Alert Impact</b>	The product of the maximum Alert Severity of alerts in the heatmap rectangle multiplied by the maximum Criticality of alerts in the heatmap rectangle. Values range from <b>0</b> - <b>10</b> , as indicated in the color gradient  bar, where <b>10</b> is the highest Alert Impact.
<b>Alert Severity</b>	<p>The maximum level of alerts in the heatmap rectangle. Values range from <b>0</b> - <b>2</b>, as indicated in the color gradient  bar, where <b>2</b> is the highest Alert Severity.</p> <ul style="list-style-type: none"> <li> Red indicates that one or more metrics have reached their alarm threshold. Metrics that have exceeded their specified ALARM LEVEL threshold have an Alert Severity value of <b>2</b>.</li> <li> Yellow indicates that one or more metrics have reached their alarm threshold. Metrics that have exceeded their specified WARNING LEVEL threshold have an Alert Severity value of <b>1</b>.</li> <li> Green indicates that no metrics have reached their alert thresholds. Metrics that have not exceeded their specified thresholds have an Alert Severity value of <b>0</b>.</li> </ul>
<b>Alert Count</b>	The total number of critical and warning alerts in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.
<b>Criticality</b>	<p>The maximum level of Criticality (rank of importance) in the heatmap rectangle. Values range from <b>1</b> to <b>5</b>, as indicated in the color gradient  bar, where <b>5</b> is the highest Criticality.</p> <p>Criticality is specified in the Service Data Model (CMDB) by your administrator. Criticality values are listed in the <b>Component Views</b> - "<a href="#">CI / Service Table</a>" display, which range from <b>A</b> to <b>E</b>, where <b>A</b> is the highest Criticality (level <b>5</b> maps to a Criticality of <b>A</b> and level <b>1</b> maps to a Criticality of <b>E</b> with equally spaced intermediate values).</p>

## Key Metrics Views

The Key Metrics (KM) feature is an entirely new way of looking at and interpreting application health and performance data.

In contrast to the traditional Alert Impact view showing your ACTIVE alerts and their impact on the overall application or service, the Key Metrics view shows how close a metric is approaching its threshold over a period of time – both before and after the alert threshold is reached.

This allows you to both proactively anticipate performance problems BEFORE the alert threshold is crossed as well analyze the circumstances that led up to error conditions AFTER you got an alert. Armed with this knowledge, you can avert disasters before they happen and resolve problems faster after they happen.

RTView does this by correlating the most valuable key metrics over multiple components within a service and displaying them in context with both real-time and historical data. This is valuable because health problems in one component may be caused by performance problems in another and only by viewing each of these metrics in context with one another over a period of time are you able to visually link the relationship between troubled components.

It is important to note that your Alert Impact heatmaps may look very different from your Key Metrics heatmaps given that KM will indicate potential threats BEFORE they show up as alerts.

Data is filtered by the \$rtvOwnerMask, \$rtvAreaMask, \$rtvGroupMask and \$rtvServiceMask values for the logged in user. For details, refer to the *RTView Enterprise Configuration Guide*. For Key Metrics definitions by technology, see ["Available KM Metrics and Alerts"](#).

## Dependencies

The KM package is dependent on the Metric Explorer package. Both must be included in your project in order for KM to work. If you are upgrading from a version previous to 1.5.0 and have not added Metric Explorer to your project, see the *RTView Enterprise User's Guide* **Upgrade Notes** section for information about including it.

Displays in this View are:

- ["Service KM Heatmap"](#): Heatmap of Key Metrics current data for one or more Services in your CMDB hierarchy.
- ["Service KM Table"](#): Table of Key Metrics current data for one or more Services.
- ["Service KM History"](#): History heatmap of Key Metrics historical data for one or more Services.
- ["Service KM History \(Alt\)"](#): History heatmap of Key Metrics historical data for one or more Services.

This section also includes:

- ["Available KM Metrics and Alerts"](#): List and descriptions of available key metrics.

## Service KM Heatmap

View Key Metrics current data for one or more Services in your CMDB hierarchy in a heatmap. The **Service KM Heatmap** provides one view of all your Services and whether they are approaching an alert condition.

The most important overview of your Services is the Alert Impact View. The Alert Impact View lets you know what is a problem NOW. The **Service KM Heatmap** gives you a proactive view of which Services might be approaching a serious problem so that you can take action before they become critical. First look at the Alert Impact View to address current issues, then move to the **Service KM Heatmap** for proactive analysis.

The colors in the display are determined by the **Threshold %** and **Quality** values. As shown in the color gradient bar , a rectangle is green when the value is close to **0** changing to yellow, orange and red as the value gets closer to **100**. Values at or over **100** are shown as red.

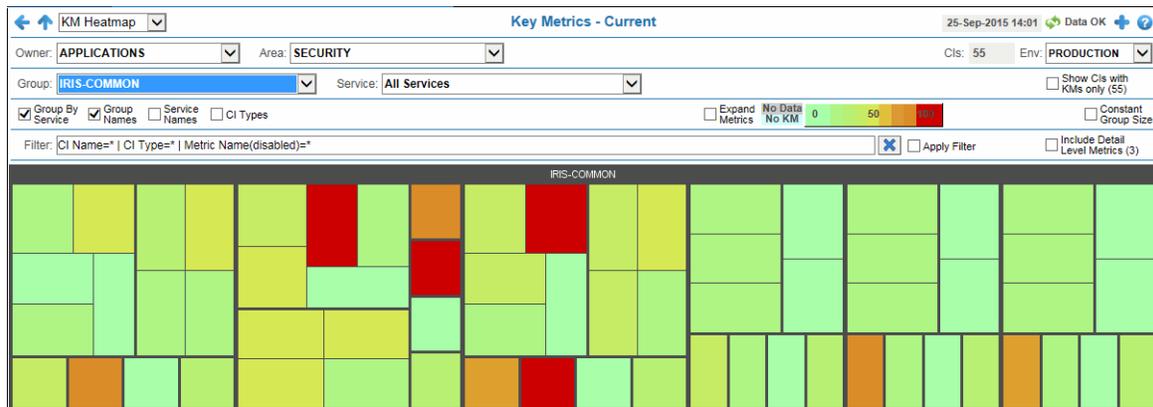
- Red indicates the value is at or over **100**.
- Yellow indicates the value is between **0** and **100**.
- Green indicates the value is close to **0**.
- Teal indicates no KMs are defined for the CI Type.
- Grey indicates KMs are defined for the CI Type but no data was returned when the metric was queried.

Select **Group By Service** to include the **Group** and **Service** labels in the display. Select **Expand Metrics** to include the **Metric Name**, **Metric Value** and **Threshold** labels in the display.

For an overview about the Key Metrics feature, see ["Key Metrics Views"](#).

For Key Metrics definitions by technology, see ["Available KM Metrics and Alerts"](#).

Use the available drop-down menus or right-click to filter data shown in the display. Drill-down and investigate by double-clicking a rectangle in the display to view details in the corresponding display.



#### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Note:** The "Up" Arrow () opens the most recently viewed display under "Multi Area Service Views". For example, if the last viewed display under **Multi Area Service Views** was **Services CI Type Summary**, then clicking opens the "Services CI Type Summary" display.

#### Filter By:

The following filtering options are typically included:

**Owner:** Choose an Owner to see metrics for Areas associated with that Owner.

**Area:** Choose an Area to see metrics for Groups associated with that Area and Owner.

**Group:** Choose a Group to see metrics for Services associated with that Group, Area and Owner.

**Service:** Choose a Service to see metrics for Environments associated with that Service, Group, Area and Owner.

**Env:** Choose an Environment to see metrics for Environments associated with that Service, Group, Area and Owner.

You can also filter KMs by **CI Name**, **CI Type** and **Metric Name**. To modify the **CI Name**, **CI Type** or **Metric Name** filter, right-click on an item in the display and select **CI Name**, **CI Type** or **Metric Name** from **Add To Filter** or **Remove From Filter**, then click **Apply Filter** . The **Filter:** field shows filter parameters and is highlighted in blue after it is applied. Click Clear  to clear the filter.

---

**Note:** The **Metric Name** filter is only editable or applied when the **Expand Metrics** checkbox is selected.

---

### Fields and Data

This display includes:

<b>Show CIs with KMs only (x)</b>	When selected, hides any CIs that do not have KMs defined. The number following the label ( <b>x</b> ) is the number of CIs with KMs defined.
<b>Group By Service</b>	When selected, includes the <b>Service Group</b> and <b>Service Name</b> in the KM data. CIs that are included in multiple Services will appear multiple times, once for each Service they are associated with.
<b>Group Names</b>	When selected, includes the <b>Group Name</b> in the display. Only available if <b>Group By Service</b> is selected.
<b>Service Names</b>	When selected, includes both the <b>Group Name</b> and <b>Service Name</b> in the display. Only available if <b>Group By Service</b> is selected.
<b>CI Types</b>	When selected, includes the CI Type in the display. If <b>Group By Service</b> is selected, this is shown in addition to the <b>Group Name</b> and <b>Service Name</b> .
<b>Expand Metrics</b>	When selected, shows one element (for example, a table row, status history row or heatmap cell) per KM per CI. When not selected, shows one element per CI with the aggregated value of all KMs for that CI. KMs are aggregated by taking the maximum <b>Threshold %</b> and the minimum <b>Quality</b> value for the CI.
	The <b>No Data No KM</b> indicates the <b>Quality</b> value for the data. If no KMs are defined for the CI Type, the <b>Quality</b> is set to <b>0</b> and the color is shown as teal. If KMs are defined for the CI Type, but no data was returned when the metric was queried, the <b>Quality</b> is set to <b>-1</b> and the color is shown as gray. If data was received for the metric, the <b>Quality</b> is set to <b>1</b> and the color is set based on the Threshold % value as described above. If the Expand Metrics checkbox is selected, this is the Quality of a single KM. If the Expand Metrics checkbox is not selected, this is the lowest Quality for all of the KMs on the CI.
	The gradient bar is the legend for the display colors, which are determined by the <b>Threshold %</b> and <b>Quality</b> values. A row is green when the value is close to <b>0</b> changing to yellow, orange and red as the value gets closer to <b>100</b> . Values at or over <b>100</b> are shown as red.
	<ul style="list-style-type: none"> <li><span style="color: red;">●</span> Red indicates the value is at or over <b>100</b>.</li> <li><span style="color: yellow;">●</span> Yellow indicates the value is between <b>0</b> and <b>100</b>.</li> <li><span style="color: green;">●</span> Green indicates the value is close to <b>0</b>.</li> <li><span style="color: teal;">●</span> Teal indicates no KMs are defined for the CI Type.</li> <li><span style="color: gray;">●</span> Grey indicates KMs are defined for the CI Type but no data was returned when the metric was queried.</li> </ul>
<b>Constant Group Size</b>	When selected, Groups are equally sized in the display. When not selected, Groups are sized according to the number of elements in the Group (a Group containing the most elements is rendered with the largest rectangle).

**Filter:** Shows the current filter parameters and is highlighted in blue when the filter is applied.  
By default, all data is shown:  
**CI Name=\* | CI Type=\* | Metric Name(disabled)=\***  
To modify the filter of KMs displayed, right-click on an item in the table and select **CI Name**, **CI Type** or **Metric Name** from **Add To Filter** or **Remove From Filter**, then click **Apply Filter** . Click Clear  to clear the filter.

Clears the filter parameters.

Applies the filter parameters.

**Include Detail Level Metrics (##)** When selected, includes **Detail Level** KMs in the display. When not selected, only includes high level KMs. The number following the label (**x**) is the number of detail level metrics available for the currently displayed KMs.

#### Mouseover

See the following details via mouseover:

**Group** The Group name. For displays showing current KM data, this column is only included if the **Group By Service** checkbox is selected. The table shows one row per Group that a CI is associated with.

**Service** The Service name. For displays showing current KM data, this column is only included if the **Group By Service** checkbox is selected. The table shows one row Service that a CI is associated with.

**CI Type** The CI Type.

**CI Name** The CI Name.

**Metric Name** The name of the metric. This is only included if the **Expand Metrics** checkbox is selected. It is the user-friendly metric name, which corresponds to a numeric column in one of the caches associated with the CI Type. To see which cache column provides data for this metric, navigate to **Architecture - "RTView KM Defs"**. In the table, look in the **CITYPE** and **SELECTOR** columns to find the row for your metric. The **CACHENAME** column lists the name of the cache containing the metric and the **METRICNAME** column contains the name of the cache column.

**Metric Value** The value of the metric. This is only included if the **Expand Metrics** checkbox is selected.

**Threshold** The **Alarm Level** value for the alert associated with the metric. This column is only included if the **Expand Metrics** checkbox is selected. To see which alert is associated with this metric, navigate to **Architecture - "RTView KM Defs"**. In the table, look in the **CITYPE** and **SELECTOR** columns to find the row for your metric. The **AlertName** column lists the name of the alert associated with the metric.  
**Note:** When looking up the alert threshold for a KM, RTView Enterprise first looks to see if there is an alert override on the alert where the **AlertIndex** exactly matches the CName (ignoring the ~ and ; delimiters). If an exact match is found, the override **Alarm Level** is used. If no exact match is found, the **Default Alarm Level** for the alert is used. Note that some alert overrides only contain a partial index and are not used for KM thresholds.

**Threshold %** The percent of the **Metric Value** against the **Threshold**. If the **Expand Metrics** checkbox is selected, this is the **Threshold %** of a single KM. If the **Expand Metrics** checkbox is not selected, this is the highest **Threshold %** for all of the KMs on the CI.

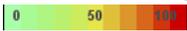
Depending on the KM, different scales are applied. By default, no scale is applied, but values are limited to **0-10000**. For memory metrics, an exponential scale is applied to the **Threshold %** so that lower values are diminished. For metrics where the alert is a low alert (an alert that executes when the value goes below the threshold), an inverse scale is applied. The colors in the KM displays are based on this value going from green when the value is close to **0** changing to yellow to orange to red as the value gets closer to **100**. Values at or over **100** are shown as red . To see which **CalcMode** is used for this metric, navigate to **Architecture - "RTView KM Defs"**. In the table, look in the **CITYPE** and **SELECTOR** columns to find the row for your metric. The **CalcMode** column lists the type of scale that is applied to the metric. If blank, no scale is applied.

<b>Quality</b>	Indicates the quality of the data. If the <b>Expand Metrics</b> checkbox is selected, the value is for a single KM on the CI. If the <b>Expand Metrics</b> checkbox is not selected, the value is for all the KMs on the CI, and shows the lowest <b>Quality</b> of those KMs. Possible values are: <b>0</b> = No KMs are defined for the CI Type (the color is shown as teal). <b>-1</b> = KMs are defined for the CI Type, but no data was returned when the metric was queried (the color is shown as gray). <b>1</b> = Data was received for the metric (the color is set based on the <b>Threshold %</b> value).
<b>Time</b>	The time stamp of the data.

## Service KM Table

View Key Metrics current data for one or more Services in your CMDB hierarchy in a table.

The **Service KM Table** shows the same information as the "[Service KM Heatmap](#)". Use this display if, for example, you prefer to sort by **Service** or **Threshold %** to identify the Service for which you want to perform proactive health analysis.

The colors of the table rows are determined by the **Threshold %** and **Quality** values. As shown in the color gradient bar , a row is green when the value is close to **0** changing to yellow, orange and red as the value gets closer to **100**. Values at or over **100** are shown as red.

- Red indicates the value is at or over **100**.
- Yellow indicates the value is between **0** and **100**.
- Green indicates the value is close to **0**.
- Teal indicates no KMs are defined for the CI Type.
- Grey indicates KMs are defined for the CI Type but no data was returned when the metric was queried.

Select **Group By Service** to include the **Group** and **Service** columns in the table. Select **Expand Metrics** to include the **Metric Name**, **Metric Value** and **Threshold** columns in the table.

---

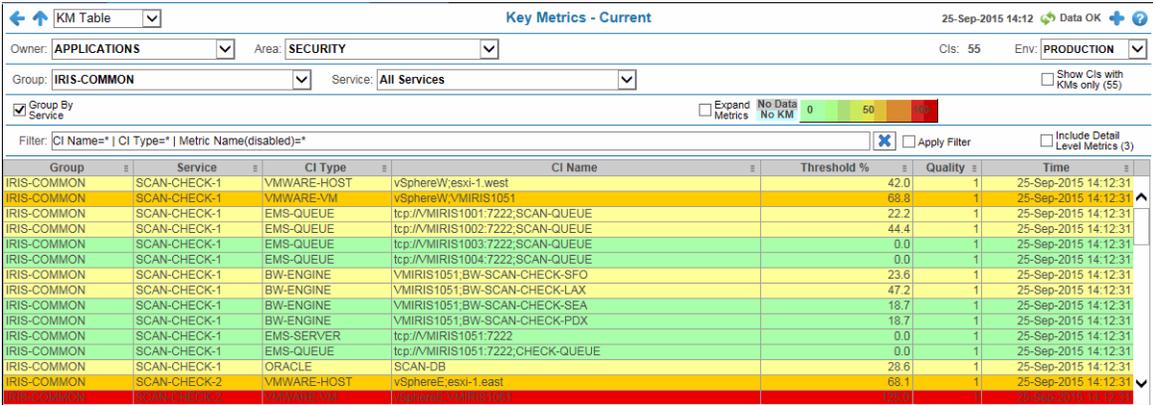
**Note:** The **CI**s label shows the number of CIs in the table. However, if the CI is associated with multiple Services it is only counted once.

---

For an overview about the Key Metrics feature, see "[Key Metrics Views](#)".

For Key Metrics definitions by technology, see "[Available KM Metrics and Alerts](#)".

Use the available drop-down menus or right-click to filter data shown in the display. Click Sort  to order column data. Drill-down and investigate by double-clicking a row in the table to view details in the corresponding display.



Group	Service	CI Type	CI Name	Threshold %	Quality	Time
IRIS-COMMON	SCAN-CHECK-1	VMWARE-HOST	vSphereW,esxi-1.west	42.0	1	25-Sep-2015 14:12:31
IRIS-COMMON	SCAN-CHECK-1	VMWARE-VM	vSphereW,VMIRIS1051	68.8	1	25-Sep-2015 14:12:31
IRIS-COMMON	SCAN-CHECK-1	EMS-QUEUE	tcp://MIRIS1001:7222:SCAN-QUEUE	22.2	1	25-Sep-2015 14:12:31
IRIS-COMMON	SCAN-CHECK-1	EMS-QUEUE	tcp://MIRIS1002:7222:SCAN-QUEUE	44.4	1	25-Sep-2015 14:12:31
IRIS-COMMON	SCAN-CHECK-1	EMS-QUEUE	tcp://MIRIS1003:7222:SCAN-QUEUE	0.0	1	25-Sep-2015 14:12:31
IRIS-COMMON	SCAN-CHECK-1	EMS-QUEUE	tcp://MIRIS1004:7222:SCAN-QUEUE	0.0	1	25-Sep-2015 14:12:31
IRIS-COMMON	SCAN-CHECK-1	BW-ENGINE	VMIRIS1051.BW-SCAN-CHECK-SFO	23.6	1	25-Sep-2015 14:12:31
IRIS-COMMON	SCAN-CHECK-1	BW-ENGINE	VMIRIS1051.BW-SCAN-CHECK-LAX	47.2	1	25-Sep-2015 14:12:31
IRIS-COMMON	SCAN-CHECK-1	BW-ENGINE	VMIRIS1051.BW-SCAN-CHECK-SEA	18.7	1	25-Sep-2015 14:12:31
IRIS-COMMON	SCAN-CHECK-1	BW-ENGINE	VMIRIS1051.BW-SCAN-CHECK-PDX	18.7	1	25-Sep-2015 14:12:31
IRIS-COMMON	SCAN-CHECK-1	EMS-SERVER	tcp://MIRIS1051:7222	0.0	1	25-Sep-2015 14:12:31
IRIS-COMMON	SCAN-CHECK-1	EMS-QUEUE	tcp://MIRIS1051:7222:CHECK-QUEUE	0.0	1	25-Sep-2015 14:12:31
IRIS-COMMON	SCAN-CHECK-1	ORACLE	SCAN-DB	28.6	1	25-Sep-2015 14:12:31
IRIS-COMMON	SCAN-CHECK-2	VMWARE-HOST	vSphereE,esxi-1.east	68.1	1	25-Sep-2015 14:12:31

**Title Bar (possible features are):**

-   Open the previous and upper display.
-  Open an instance of this display in a new window.
-  Open the online help page for this display.
- open commonly accessed displays.

 **Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Note:** The "Up" Arrow () opens the most recently viewed display under "Multi Area Service Views". For example, if the last viewed display under **Multi Area Service Views** was **Services CI Type Summary**, then clicking  opens the "Services CI Type Summary" display.

#### Filter By:

The following filtering options are typically included:

**Owner:** Choose an Owner to see metrics for Areas associated with that Owner.

**Area:** Choose an Area to see metrics for Groups associated with that Area and Owner.

**Group:** Choose a Group to see metrics for Services associated with that Group, Area and Owner.

**Service:** Choose a Service to see metrics for Environments associated with that Service, Group, Area and Owner.

**Env:** Choose an Environment to see metrics for Environments associated with that Service, Group, Area and Owner.

You can also filter KMs by **CI Name**, **CI Type** and **Metric Name**. To modify the **CI Name**, **CI Type** or **Metric Name** filter, right-click on an item in the display and select **CI Name**, **CI Type** or **Metric Name** from **Add To Filter** or **Remove From Filter**, then click **Apply Filter** . The **Filter:** field shows filter parameters and is highlighted in blue after it is applied. Click Clear  to clear the filter.

**Note:** The **Metric Name** filter is only editable or applied when the **Expand Metrics** checkbox is selected.

**Fields and Data**

This display includes:

**Show CIs with KMs only (x)**

When selected, hides any CIs that do not have KMs defined. The number following the label (**x**) is the number of CIs with KMs defined.

**Group By Service**

When selected, includes the **Service Group** and **Service Name** in the KM data. CIs that are included in multiple Services will appear multiple times, once for each Service they are associated with.

**Expand Metrics**

When selected, shows one element (for example, a table row, status history row or heatmap cell) per KM per CI. When not selected, shows one element per CI with the aggregated value of all KMs for that CI. KMs are aggregated by taking the maximum **Threshold %** and the minimum **Quality** value for the CI.



The **No Data No KM** indicates the **Quality** value for the data. If no KMs are defined for the CI Type, the **Quality** is set to **0** and the color is shown as teal. If KMs are defined for the CI Type, but no data was returned when the metric was queried, the **Quality** is set to **-1** and the color is shown as gray. If data was received for the metric, the **Quality** is set to **1** and the color is set based on the **Threshold %** value as described above. If the **Expand Metrics** checkbox is selected, this is the Quality of a single KM. If the **Expand Metrics** checkbox is not selected, this is the lowest Quality for all of the KMs on the CI.



The gradient bar is the legend for the table row colors, which are determined by the **Threshold %** and **Quality** values. A row is green when the value is close to **0** changing to yellow, orange and red as the value gets closer to **100**. Values at or over **100** are shown as red.

- Red indicates the value is at or over **100**.
- Yellow indicates the value is between **0** and **100**.
- Green indicates the value is close to **0**.
- Teal indicates no KMs are defined for the CI Type.
- Grey indicates KMs are defined for the CI Type but no data was returned when the metric was queried.

**Filter:**

Shows the current filter parameters and is highlighted in blue when the filter is applied.

By default, all data is shown:

**CI Name=\* | CI Type=\* | Metric Name(disabled)=\***

To modify the filter of KMs displayed, right-click on an item in the table and select **CI Name**, **CI Type** or **Metric Name** from **Add To Filter** or **Remove From Filter**, then click **Apply Filter** . Click **Clear**  to clear the filter.



Clears the filter parameters.



Applies the filter parameters.

**Include Detail Level Metrics (##)**

When selected, includes **Detail Level** KMs in the display. When not selected, only includes high level KMs. The number following the label (**x**) is the number of detail level metrics available for the currently displayed KMs.

**Group**

The Group name. For displays showing current KM data, this column is only included if the **Group By Service** checkbox is selected. The table shows one row per Group that a CI is associated with.

**Service**

The Service name. For displays showing current KM data, this column is only included if the **Group By Service** checkbox is selected. The table shows one row Service that a CI is associated with.

**CI Type**

The CI Type.

**CI Name**

The CI Name.

<b>Metric Name</b>	The name of the metric. This column is only included if the <b>Expand Metrics</b> checkbox is selected. It is the user-friendly metric name, which corresponds to a numeric column in one of the caches associated with the CI Type. To see which cache column provides data for this metric, navigate to <b>Architecture - "RTView KM Defs"</b> . In the table, look in the <b>CITYPE</b> and <b>SELECTOR</b> columns to find the row for your metric. The <b>CACHENAME</b> column lists the name of the cache containing the metric and the <b>METRICNAME</b> column contains the name of the cache column.
<b>Metric Value</b>	The value of the metric. This column is only included if the <b>Expand Metrics</b> checkbox is selected.
<b>Threshold</b>	The <b>Alarm Level</b> value for the alert associated with the metric. This column is only included if the <b>Expand Metrics</b> checkbox is selected. To see which alert is associated with this metric, navigate to <b>Architecture - "RTView KM Defs"</b> . In the table, look in the <b>CITYPE</b> and <b>SELECTOR</b> columns to find the row for your metric. The <b>AlertName</b> column lists the name of the alert associated with the metric. <b>Note:</b> When looking up the alert threshold for a KM, RTView Enterprise first looks to see if there is an alert override on the alert where the <b>AlertIndex</b> exactly matches the CIName (ignoring the <b>~</b> and <b>;</b> delimiters). If an exact match is found, the override <b>Alarm Level</b> is used. If no exact match is found, the <b>Default Alarm Level</b> for the alert is used. Note that some alert overrides only contain a partial index and are not used for KM thresholds.
<b>Threshold %</b>	The percent of the <b>Metric Value</b> against the <b>Threshold</b> . If the <b>Expand Metrics</b> checkbox is selected, this is the <b>Threshold %</b> of a single KM. If the <b>Expand Metrics</b> checkbox is not selected, this is the highest <b>Threshold %</b> for all of the KMs on the CI.  Depending on the KM, different scales are applied. By default, no scale is applied, but values are limited to <b>0-10000</b> . For memory metrics, an exponential scale is applied to the <b>Threshold %</b> so that lower values are diminished. For metrics where the alert is a low alert (an alert that executes when the value goes below the threshold), an inverse scale is applied. The colors in the KM displays are based on this value going from green when the value is close to <b>0</b> changing to yellow to orange to red as the value gets closer to <b>100</b> . Values at or over <b>100</b> are shown as red  . To see which <b>CalcMode</b> is used for this metric, navigate to <b>Architecture - "RTView KM Defs"</b> . In the table, look in the <b>CITYPE</b> and <b>SELECTOR</b> columns to find the row for your metric. The <b>CalcMode</b> column lists the type of scale that is applied to the metric. If blank, no scale is applied.
<b>Quality</b>	Indicates the quality of the data. If the <b>Expand Metrics</b> checkbox is selected, the value is for a single KM on the CI. If the <b>Expand Metrics</b> checkbox is not selected, the value is for all the KMs on the CI, and shows the lowest <b>Quality</b> of those KMs. Possible values are: <b>0</b> = No KMs are defined for the CI Type (the color is shown as teal). <b>-1</b> = KMs are defined for the CI Type, but no data was returned when the metric was queried (the color is shown as gray). <b>1</b> = Data was received for the metric (the color is set based on the <b>Threshold %</b> value).
<b>Time</b>	The time stamp of the data.

## Service KM History

View history heatmap of Key Metric data, over time, for a selected Group and Service.

This is the most important view for analyzing the correlation between a variety of Key Metrics over time that are related to a Service. You would navigate to this view if:

- you have identified a Service in the Alert Impact View that is having degradation right now. You can select the Service and navigate to the **Service KM History** display to determine if there are various factors causing the degradation.
- you have looked at the **"Service KM Heatmap"** or the **"Service KM Table"** and identified a Service that is about to become degraded. You can navigate to the **Service KM History** display to proactively analyze the Service before issues arise.

Each row in the history heatmap represents a different CI, unless the **Expand Metrics** checkbox is selected, in which case it represents a metric on a CI. The row color shows the **Threshold %** and **Quality** values.

The **Threshold %** value is rounded up to the closest **10** unless the **Quality** is less than **1**, in which case it shows the **Quality**. As shown in the color gradient bar , the color is green when the value is close to **0** changing to yellow, orange and red as the value gets closer to **100**. Values at or over **100** are shown as red.

- Red indicates the value is at or over **100**.
- Yellow indicates the value is between **0** and **100**.
- Green indicates the value is close to **0**.
- Teal indicates no KMs are defined for the CI Type.
- Grey indicates KMs are defined for the CI Type but no data was returned when the metric was queried.

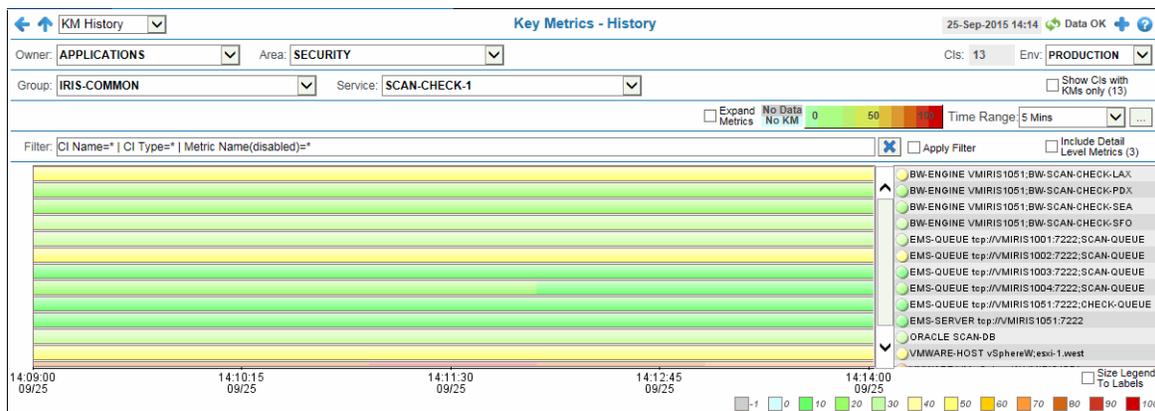
CIs associated with multiple Services are shown once for the first **Group** and **Service** they were associated with, and labeled **+ x more**, where **x** is the number of additional **Groups** and **Services** the CI is associated with.

Select **Expand Metrics** to show each Key Metric in its own row and include the **Metric Name**, **Metric Value** and **Threshold** labels in the mouseover popup window.

For an overview about the Key Metrics feature, see ["Key Metrics Views"](#).

For Key Metrics definitions by technology, see ["Available KM Metrics and Alerts"](#)

Use the available drop-down menus or right-click to filter data shown in the display. Drill-down and investigate by double-clicking a row to view details in the corresponding display.



#### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Note:** The "Up" Arrow () opens the most recently viewed display under ["Multi Area Service Views"](#). For example, if the last viewed display under **Multi Area Service Views** was **Services CI Type Summary**, then clicking opens the ["Services CI Type Summary"](#) display.

**Filter By:**

The following filtering options are typically included:

**Owner:** Choose an Owner to see metrics for Areas associated with that Owner.

**Area:** Choose an Area to see metrics for Groups associated with that Area and Owner.

**Group:** Choose a Group to see metrics for Services associated with that Group, Area and Owner.

**Service:** Choose a Service to see metrics for Environments associated with that Service, Group, Area and Owner.

**Env:** Choose an Environment to see metrics for Environments associated with that Service, Group, Area and Owner.

You can also filter KMs by CI Name, CI Type and Metric Name. To modify the **CI Name**, **CI Type** or **Metric Name** filter, right-click on an item and select **CI Name**, **CI Type** or **Metric Name** from **Add To Filter** or **Remove From Filter**, then click **Apply Filter** . The **Filter:** field shows filter parameters and is highlighted in blue after it is applied. Click **Clear**  to clear the filter.

---

**Note:** The **Metric Name** filter is only editable or applied when the **Expand Metrics** checkbox is selected.

---

**Fields and Data**

This display includes:

**Show CIs with KMs only (x)**

When selected, hides any CIs that do not have KMs defined. The number following the label (x) is the number of CIs with KMs defined.

**Expand Metrics**

When selected, shows one element (for example, a table row, status history row or heatmap cell) per KM per CI. When not selected, shows one element per CI with the aggregated value of all KMs for that CI. KMs are aggregated by taking the maximum **Threshold %** and the minimum **Quality** value for the CI.



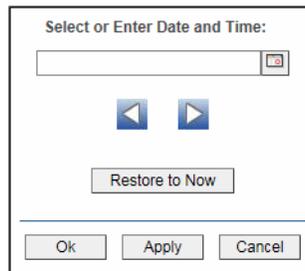
The **No Data No KM** is the legend for the display colors if the **Quality** value for the data is less than **1**. If no KMs are defined for the **CI Type**, the **Quality** is set to **0** and the color is shown as teal. If KMs are defined for the CI Type, but no data was returned when the metric was queried, the **Quality** is set to **-1** and the color is shown as gray. If data was received for the metric, the **Quality** is set to **1** and the color is set based on the **Threshold %** value as described above. If the **Expand Metrics** checkbox is selected, this is the **Quality** of a single KM. If the **Expand Metrics** checkbox is not selected, this is the lowest **Quality** for all of the KMs on the CI.



The gradient bar is the legend for the display colors, which are determined by the **Threshold %** and **Quality** values. A row is green when the value is close to **0** changing to yellow, orange and red as the value gets closer to **100**. Values at or over **100** are shown as red.

-  Red indicates the value is at or over **100**.
-  Yellow indicates the value is between **0** and **100**.
-  Green indicates the value is close to **0**.
-  Teal indicates no KMs are defined for the CI Type.
-  Grey indicates KMs are defined for the CI Type but no data was returned when the metric was queried.

**Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

**Note:** To limit the memory used by the KM history displays, the available time ranges are limited by the number of CIs in the selected **Group** and **Service**. This limit can be modified using the **\$rtvKMHistroyRowLimit** substitution. The **\$rtvKMHistroyRowLimit** substitution sets the maximum number of rows that can be queried by a history display and this number is used to determine the available time ranges. The default value is **35000**. To change the limit (and the maximum amount of memory used by KM history display), set the following property to a different value: **sl.rtvview.sub=\$rtvKMHistroyRowLimit:35000**.

**Filter:** Shows the current filter parameters and is highlighted in blue when the filter is applied.

By default, all data is shown:

**CI Name=\* | CI Type=\* | Metric Name(disabled)=\***

To modify the filter of KMs displayed, right-click on an item in the table and select **CI Name**, **CI Type** or **Metric Name** from **Add To Filter** or **Remove From Filter**, then click **Apply Filter** . Click **Clear**  to clear the filter.



Clears the filter parameters.



Applies the filter parameters.

**Include Detail Level Metrics (##)** When selected, includes **Detail Level** KMs in the display. When not selected, only includes high level KMs. The number following the label (**x**) is the number of detail level metrics available for the currently displayed KMs.

#### Mouseover

See the following details via mouseover:

**Group** The **Group** name. For CIs that are associated with multiple Groups, the name of the first **Group** the CI was associated with is shown and labeled **+ x more**, where **x** is the number of additional **Groups** the CI is associated with.

**Service** The **Service** name. For CIs that are associated with multiple Services, the name of the first **Service** the CI was associated with is shown and labeled **+ x more**, where **x** is the number of additional **Services** the CI is associated with.

**CI Type** The CI Type.

**CI Name** The CI Name.

<b>Metric Name</b>	The name of the metric. This is only included if the <b>Expand Metrics</b> checkbox is selected. It is the user-friendly metric name, which corresponds to a numeric column in one of the caches associated with the <b>CI Type</b> . To see which cache column provides data for this metric, navigate to <b>Architecture - "RTView KM Defs"</b> . In the table, look in the <b>CITYPE</b> and <b>SELECTOR</b> columns to find the row for your metric. The <b>CACHENAME</b> column lists the name of the cache containing the metric and the <b>METRICNAME</b> column contains the name of the cache column.
<b>Metric Value</b>	The value of the metric. This is only included if the <b>Expand Metrics</b> checkbox is selected.
<b>Threshold</b>	The <b>Alarm Level</b> value for the alert associated with the metric. This column is only included if the <b>Expand Metrics</b> checkbox is selected. To see which alert is associated with this metric, navigate to <b>Architecture - "RTView KM Defs"</b> . In the table, look in the <b>CITYPE</b> and <b>SELECTOR</b> columns to find the row for your metric. The <b>AlertName</b> column lists the name of the alert associated with the metric. <b>Note:</b> When looking up the alert threshold for a KM, RTView Enterprise first looks to see if there is an alert override on the alert where the <b>AlertIndex</b> exactly matches the CIName (ignoring the ~ and ; delimiters). If an exact match is found, the override <b>Alarm Level</b> is used. If no exact match is found, the <b>Default Alarm Level</b> for the alert is used. Note that some alert overrides only contain a partial index and are not used for KM thresholds.
<b>Threshold %</b>	The percent of the <b>Metric Value</b> against the <b>Threshold</b> . If the <b>Expand Metrics</b> checkbox is selected, this is the <b>Threshold %</b> of a single KM. If the <b>Expand Metrics</b> checkbox is not selected, this is the highest <b>Threshold %</b> for all of the KMs on the CI.  Depending on the KM, different scales are applied. By default, no scale is applied, but values are limited to <b>0-10000</b> . For memory metrics, an exponential scale is applied to the <b>Threshold %</b> so that lower values are diminished. For metrics where the alert is a low alert (an alert that executes when the value goes below the threshold), an inverse scale is applied. The colors in the KM displays are based on this value going from green when the value is close to <b>0</b> changing to yellow to orange to red as the value gets closer to <b>100</b> . Values at or over <b>100</b> are shown as red  . To see which <b>CalcMode</b> is used for this metric, navigate to <b>Architecture - "RTView KM Defs"</b> . In the table, look in the <b>CITYPE</b> and <b>SELECTOR</b> columns to find the row for your metric. The <b>CalcMode</b> column lists the type of scale that is applied to the metric. If blank, no scale is applied.
<b>Quality</b>	Indicates the quality of the data. If the <b>Expand Metrics</b> checkbox is selected, the value is for a single KM on the CI. If the <b>Expand Metrics</b> checkbox is not selected, the value is for all the KMs on the CI, and shows the lowest <b>Quality</b> of those KMs. Possible values are: <b>0</b> = No KMs are defined for the CI Type (the color is shown as teal). <b>-1</b> = KMs are defined for the CI Type, but no data was returned when the metric was queried (the color is shown as gray). <b>1</b> = Data was received for the metric (the color is set based on the <b>Threshold %</b> value).
<b>Time</b>	The time stamp of the data.
<b>Size Legend To Labels</b>	When selected, the width of the legend is set to the widest label. When not selected, the width of the legend is set to 20% of the available space and labels that are too wide are clipped.

## Service KM History (Alt)

View history heatmap of Key Metric data, over time, for a selected Group and Service. This display shows the same data as the **"Service KM History"** display but contains fewer labels. Each row in the history heatmap represents a different CI, unless the **Expand Metrics** checkbox is selected, in which case it represents a metric on a CI. The row color shows the **Threshold %** and **Quality** values.

As shown in the color gradient bar , the color is green when the value is close to **0** changing to yellow, orange and red as the value gets closer to **100**. Values at or over **100** are shown as red.

- Red indicates the value is at or over **100**.
- Yellow indicates the value is between **0** and **100**.
- Green indicates the value is close to **0**.
- Teal indicates no KMs are defined for the CI Type.
- Grey indicates KMs are defined for the CI Type but no data was returned when the metric was queried.

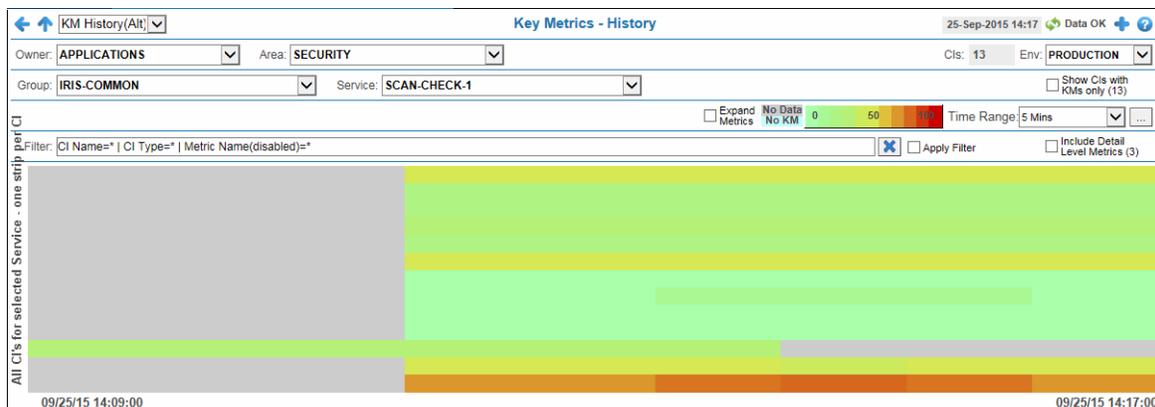
CIs associated with multiple Services are shown once for the first **Group** and **Service** they were associated with, and labeled **+ x more**, where **x** is the number of additional **Groups** and **Services** the CI is associated with.

Select **Expand Metrics** to show each Key Metric in its own row and include the **Metric Name**, **Metric Value** and **Threshold** labels in the mouseover popup window.

For an overview about the Key Metrics feature, see "[Key Metrics Views](#)".

For Key Metrics definitions by technology, see "[Available KM Metrics and Alerts](#)".

Use the available drop-down menus or right-click to filter data shown in the display. Drill-down and investigate by double-clicking a row to view details in the corresponding display.



#### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Note:** The "Up" Arrow () opens the most recently viewed display under "[Multi Area Service Views](#)". For example, if the last viewed display under **Multi Area Service Views** was **Services CI Type Summary**, then clicking opens the "[Services CI Type Summary](#)" display.

**Filter By:**

The following filtering options are typically included:

**Owner:** Choose an Owner to see metrics for Areas associated with that Owner.

**Area:** Choose an Area to see metrics for Groups associated with that Area and Owner.

**Group:** Choose a Group to see metrics for Services associated with that Group, Area and Owner.

**Service:** Choose a Service to see metrics for Environments associated with that Service, Group, Area and Owner.

**Env:** Choose an Environment to see metrics for Environments associated with that Service, Group, Area and Owner.

You can also filter KMs by CI Name, CI Type and Metric Name. To modify the **CI Name**, **CI Type** or **Metric Name** filter, right-click on an item and select **CI Name**, **CI Type** or **Metric Name** from **Add To Filter** or **Remove From Filter**, then click **Apply Filter** . The **Filter:** field shows filter parameters and is highlighted in blue after it is applied. Click **Clear**  to clear the filter.

---

**Note:** The **Metric Name** filter is only editable or applied when the **Expand Metrics** checkbox is selected.

---

**Fields and Data**

This display includes:

**Show CIs with KMs only (x)**

When selected, hides any CIs that do not have KMs defined. The number following the label (x) is the number of CIs with KMs defined.

**Expand Metrics**

When selected, shows one element (for example, a table row, status history row or heatmap cell) per KM per CI. When not selected, shows one element per CI with the aggregated value of all KMs for that CI. KMs are aggregated by taking the maximum **Threshold %** and the minimum **Quality** value for the CI.



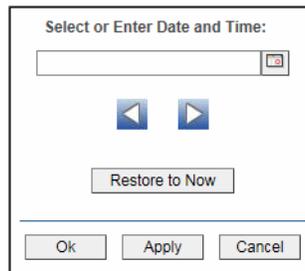
The **No Data No KM** is the legend for the display colors if the **Quality** value for the data is less than **1**. If no KMs are defined for the **CI Type**, the **Quality** is set to **0** and the color is shown as teal. If KMs are defined for the CI Type, but no data was returned when the metric was queried, the **Quality** is set to **-1** and the color is shown as gray. If data was received for the metric, the **Quality** is set to **1** and the color is set based on the **Threshold %** value as described above. If the **Expand Metrics** checkbox is selected, this is the **Quality** of a single KM. If the **Expand Metrics** checkbox is not selected, this is the lowest **Quality** for all of the KMs on the CI.



The gradient bar is the legend for the display colors, which are determined by the **Threshold %** and **Quality** values. A row is green when the value is close to **0** changing to yellow, orange and red as the value gets closer to **100**. Values at or over **100** are shown as red.

-  Red indicates the value is at or over **100**.
-  Yellow indicates the value is between **0** and **100**.
-  Green indicates the value is close to **0**.
-  Teal indicates no KMs are defined for the CI Type.
-  Grey indicates KMs are defined for the CI Type but no data was returned when the metric was queried.

**Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

**Note:** To limit the memory used by the KM history displays, the available time ranges are limited by the number of CIs in the selected **Group** and **Service**. This limit can be modified using the **\$rtvKMHistroyRowLimit** substitution. The **\$rtvKMHistroyRowLimit** substitution sets the maximum number of rows that can be queried by a history display and this number is used to determine the available time ranges. The default value is **35000**. To change the limit (and the maximum amount of memory used by KM history display), set the following property to a different value: **sl.rtvview.sub=\$rtvKMHistroyRowLimit:35000**.

**Filter:** Shows the current filter parameters and is highlighted in blue when the filter is applied.

By default, all data is shown:

**CI Name=\* | CI Type=\* | Metric Name(disabled)=\***

To modify the filter of KMs displayed, right-click on an item in the table and select **CI Name**, **CI Type** or **Metric Name** from **Add To Filter** or **Remove From Filter**, then click **Apply Filter** . Click Clear  to clear the filter.

 Clears the filter parameters.

 Applies the filter parameters.

**Include Detail Level Metrics (##)** When selected, includes **Detail Level** KMs in the display. When not selected, only includes high level KMs. The number following the label (**x**) is the number of detail level metrics available for the currently displayed KMs.

#### Mouseover

See the following details via mouseover:

**Group** The **Group** name. For CIs that are associated with multiple Groups, the name of the first **Group** the CI was associated with is shown and labeled **+ x more**, where **x** is the number of additional **Groups** the CI is associated with.

**Service** The **Service** name. For CIs that are associated with multiple Services, the name of the first **Service** the CI was associated with is shown and labeled **+ x more**, where **x** is the number of additional **Services** the CI is associated with.

**CI Type** The CI Type.

**CI Name** The CI Name.

<b>Metric Name</b>	The name of the metric. This is only included if the <b>Expand Metrics</b> checkbox is selected. It is the user-friendly metric name, which corresponds to a numeric column in one of the caches associated with the <b>CI Type</b> . To see which cache column provides data for this metric, navigate to <b>Architecture - "RTView KM Defs"</b> . In the table, look in the <b>CITYPE</b> and <b>SELECTOR</b> columns to find the row for your metric. The <b>CACHENAME</b> column lists the name of the cache containing the metric and the <b>METRICNAME</b> column contains the name of the cache column.
<b>Metric Value</b>	The value of the metric. This is only included if the <b>Expand Metrics</b> checkbox is selected.
<b>Threshold</b>	The <b>Alarm Level</b> value for the alert associated with the metric. This column is only included if the <b>Expand Metrics</b> checkbox is selected. To see which alert is associated with this metric, navigate to <b>Architecture - "RTView KM Defs"</b> . In the table, look in the <b>CITYPE</b> and <b>SELECTOR</b> columns to find the row for your metric. The <b>AlertName</b> column lists the name of the alert associated with the metric. <b>Note:</b> When looking up the alert threshold for a KM, RTView Enterprise first looks to see if there is an alert override on the alert where the <b>AlertIndex</b> exactly matches the CName (ignoring the ~ and ; delimiters). If an exact match is found, the override <b>Alarm Level</b> is used. If no exact match is found, the <b>Default Alarm Level</b> for the alert is used. Note that some alert overrides only contain a partial index and are not used for KM thresholds.
<b>Threshold %</b>	The percent of the <b>Metric Value</b> against the <b>Threshold</b> . If the <b>Expand Metrics</b> checkbox is selected, this is the <b>Threshold %</b> of a single KM. If the <b>Expand Metrics</b> checkbox is not selected, this is the highest <b>Threshold %</b> for all of the KMs on the CI.  Depending on the KM, different scales are applied. By default, no scale is applied, but values are limited to <b>0-10000</b> . For memory metrics, an exponential scale is applied to the <b>Threshold %</b> so that lower values are diminished. For metrics where the alert is a low alert (an alert that executes when the value goes below the threshold), an inverse scale is applied. The colors in the KM displays are based on this value going from green when the value is close to <b>0</b> changing to yellow to orange to red as the value gets closer to <b>100</b> . Values at or over <b>100</b> are shown as red  . To see which <b>CalcMode</b> is used for this metric, navigate to <b>Architecture - "RTView KM Defs"</b> . In the table, look in the <b>CITYPE</b> and <b>SELECTOR</b> columns to find the row for your metric. The <b>CalcMode</b> column lists the type of scale that is applied to the metric. If blank, no scale is applied.
<b>Quality</b>	Indicates the quality of the data. If the <b>Expand Metrics</b> checkbox is selected, the value is for a single KM on the CI. If the <b>Expand Metrics</b> checkbox is not selected, the value is for all the KMs on the CI, and shows the lowest <b>Quality</b> of those KMs. Possible values are:  <b>0</b> = No KMs are defined for the CI Type (the color is shown as teal). <b>-1</b> = KMs are defined for the CI Type, but no data was returned when the metric was queried (the color is shown as gray). <b>1</b> = Data was received for the metric (the color is set based on the <b>Threshold %</b> value).
<b>Time</b>	The time stamp of the data.

## Available KM Metrics and Alerts

This section lists available Key Metrics and their associated alerts.

- "Amazon AWS"
- "Host Agent"
- "IBM DB2"
- "IBM MQ"
- "IBM WebSphere"
- "JBoss"
- "Kafka"
- "Oracle Coherence"
- "Oracle Database"
- "Oracle WebLogic"
- "RTViewManager"
- "RTVRULES"
- "Solace"
- "TIBCO ActiveMatrix"
- "TIBCO ActiveSpaces"
- "TIBCO BusinessEvents"
- "TIBCO BusinessWorks (Version 5) Monitor"
- "TIBCO BusinessWorks (Version 6) Monitor"
- "TIBCO EMS"
- "UX"

### Amazon AWS

The following KMs are available with the Solution Package for Amazon AWS. The level of the Key Metric is **0** except where noted. Level **0** KMs are always shown in displays. Level **1** KMs are displayed if the **Include Detail Level Metrics** checkbox is selected.)

CI Type	Cache	Selector	Metric / Alert
ACW	AwsEc2InstanceStats	Instance CPU Usage	CPUUtilization / AcwInstanceCpuHigh

### Host Agent

The following KMs are available with the Solution Package for RTView Host Agent. The level of the Key Metric is **0** except where noted. Level **0** KMs are always shown in displays. Level **1** KMs are displayed if the **Include Detail Level Metrics** checkbox is selected.)

CI Type	Cache	Selector	Metric / Alert
HOST	HostStats	% CPU Utilization	usedPerCentCpu / HostCpuPercentHigh
HOST	HostStats	% Memory Used	MemUsedPerCent / HostMemoryUsedHigh

### IBM DB2

The following KMs are available with the Solution Package for IBM DB2. The level of the Key Metric is **0** except where noted. Level **0** KMs are always shown in displays. Level **1** KMs are displayed if the **Include Detail Level Metrics** checkbox is selected.)

CI Type	Cache	Selector	Metric / Alert
DB2	Db2ResponseTime	Response Time	ResponseTimeMilliSec / Db2ResponseTimeHigh

## IBM MQ

The following KMs are available with the Solution Package for IBM WebSphere MQ. The level of the Key Metric is **0** except where noted. Level **0** KMs are always shown in displays. Level **1** KMs are displayed if the **Include Detail Level Metrics** checkbox is selected.)

CI Type	Cache	Selector	Metric / Alert
<b>MQ-BROKER</b>	MqBrokers	Queue Depth	<b>Current queue depth / MqBrokerQueueDepthHigh</b>
<b>MQ-QUEUE</b>	MqQueues	Queue Depth	<b>Current queue depth / MqQueueDepthHigh</b>

## IBM WebSphere

The following KMs are available with the Solution Package for IBM WebSphere. The level of the Key Metric is **0** except where noted. Level **0** KMs are always shown in displays. Level **1** KMs are displayed if the **Include Detail Level Metrics** checkbox is selected.)

CI Type	Cache	Selector	Metric / Alert
<b>WAS</b>	WasServerStats	Live Session Count	<b>LiveCount / WasLiveSessionCountHigh</b>
<b>WAS</b>	WasServerStats	WAS CPU %	<b>ProcessCpuUsage / WasJvmCpuHigh</b>
<b>WAS</b>	WasServerStats	Memory Used %	<b>usedMemoryPercent / WasMemoryUsedPercentHigh</b>
<b>WAS-APP</b>	WasServletTotalsByApp	Response Time	<b>responseTime / WasServletResponseTimeHigh</b>
<b>WAS-APP</b>	WasServletTotalsByApp	Requests / sec	<b>DeltatotalRequests / WasServletRequestRateHigh</b>

## JBoss

The following KMs are available with the Solution Package for JBoss. The level of the Key Metric is **0** except where noted. Level **0** KMs are always shown in displays. Level **1** KMs are displayed if the **Include Detail Level Metrics** checkbox is selected.)

CI Type	Cache	Selector	Metric / Alert
<b>JBOSS-APP</b>	JbossDeployments	Active Sessions	<b>activeSessions / JbossAppActiveSessionsHigh</b>
<b>JBOSS-SERVER</b>	JbossServerStats	% Process CPU	<b>ProcessCpuLoadPercent / JbossServerProcessCpuLoadHigh</b>
<b>JBOSS-SERVER</b>	JbossDeploymentTotals	Active Sessions	<b>activeSessions / JbossServerActiveSessionsHigh</b> The level of this Key Metric is <b>1</b> . (Level <b>0</b> KMs are always displayed. Level <b>1</b> KMs are displayed if the <b>Include Detail Level Metrics</b> checkbox is checked.)

## Kafka

The following KMs are available with the Solution Package for Kafka. The level of the Key Metric is **0** except where noted. Level **0** KMs are always shown in displays. Level **1** KMs are displayed if the **Include Detail Level Metrics** checkbox is selected.

CI Type	Cache	Selector	Metric / Alert
<b>KAFKA-BROKER</b>	KafkaServerJvm	CpuPercent	<b>CpuPercent / KafkaBrokerCpuPercentHigh</b>
<b>KAFKA-BROKER</b>	KafkaServerJvm	MemoryUsedPercent	<b>MemoryUsedPercent / KafkaBrokerMemoryUsedPercentHigh</b>
<b>KAFKA-CONSUMER</b>	KafkaConsumer	Consumer Lag	<b>records-lag-max / KafkaConsumerLagHigh</b>
<b>KAFKA-CONSUMER</b>	KafkaConsumer	CpuPercent	<b>CpuPercent / KafkaConsumerCpuPercentHigh</b>
<b>KAFKA-CONSUMER</b>	KafkaConsumer	MemoryUsedPercent	<b>MemoryUsedPercent / KafkaConsumerMemoryUsedPercentHigh</b>
<b>KAFKA-PRODUCER</b>	KafkaProducer	CpuPercent	<b>CpuPercent / KafkaProducerCpuPercentHigh</b>
<b>KAFKA-PRODUCER</b>	KafkaProducer	IO Wait Time	<b>io-wait-time-millis-avg / KafkaProducerIoWaitTimeMSHigh</b>
<b>KAFKA-PRODUCER</b>	KafkaProducer	MemoryUsedPercent	<b>MemoryUsedPercent / KafkaProducerMemoryUsedPercentHigh</b>
<b>KAFKA-ZOOKEEPER</b>	KafkaZookeeper	Avg Request Latency	<b>AvgRequestLatency / KafkaZookeeperAvgLatencyHigh</b>

## Oracle Coherence

The following KMs are available with the Solution Package for Oracle Coherence. The level of the Key Metric is **0** except where noted. Level **0** KMs are always shown in displays. Level **1** KMs are displayed if the **Include Detail Level Metrics** checkbox is selected.) For alert definitions, see *RTView® Oracle® Coherence Monitor User's Guide*.

CI Type	Cache	Selector	Metric / Alert
<b>OC-CACHE</b>	OcCacheTotals	Rate Cache Misses	<b>RateCacheMisses / OcCacheRateCacheMissesHigh</b> This metric is the rate of cache misses against a given tier of a given cache for a given service in a given (Coherence) cluster. The tier can be front, where appropriate, or back. Caches and services are named, and (Coherence) clusters are represented by their named monitoring connection.

<b>OC-CACHE</b>	OcCacheTotals	Rate Store Reads	<p><b>RateStoreReads / OcCacheRateStoreReadsHigh</b></p> <p>The level of this Key Metric is <b>1</b>. (Level <b>0</b> KMs are always displayed. Level <b>1</b> KMs are displayed if the <b>Include Detail Level Metrics</b> checkbox is checked.)</p> <p>This metric is the rate of store reads (load operations) against a given tier of a given cache for a given service in a given (Coherence) cluster. The tier can be front, where appropriate, or back. Caches and services are named, and (Coherence) clusters are represented by their named monitoring connection.</p>
<b>OC-CACHE</b>	OcCacheTotals	Rate Store Writes	<p><b>RateStoreWrites / OcCacheRateStoreWritesHigh</b></p> <p>The level of this Key Metric is <b>1</b>. (Level <b>0</b> KMs are always displayed. Level <b>1</b> KMs are displayed if the <b>Include Detail Level Metrics</b> checkbox is checked.)</p> <p>This metric is the rate of store writes (store and erase operations) against a given tier of a given cache for a given service in a given (Coherence) cluster. The tier can be front, where appropriate, or back. Caches and services are named, and (Coherence) clusters are represented by their named monitoring connection.</p>
<b>OC-CACHE</b>	OcCacheTotals	Queue Size	<p><b>QueueSizePos / OcCacheQueueSizeHigh</b></p> <p>This metric is the cache send queue size for a given tier of a given cache for a given service in a given (Coherence) cluster. The tier can be front, where appropriate, or back. Caches and services are named, and (Coherence) clusters are represented by their named monitoring connection.</p>
<b>OC-CACHE</b>	OcCacheTotals	Rate Cache Puts	<p><b>RateTotalPuts / OcCacheRateTotalPutsHigh</b></p> <p>This metric is the rate of cache puts against a given tier of a given cache for a given service in a given (Coherence) cluster. The tier can be front, where appropriate, or back. Caches and services are named, and (Coherence) clusters are represented by their named monitoring connection.</p>
<b>OC-CACHE</b>	OcCacheTotals	Rate Cache Gets	<p><b>RateTotalGets / OcCacheRateTotalGetsHigh</b></p> <p>This metric is the rate of cache gets against a given tier of a given cache for a given service in a given (Coherence) cluster. The tier can be front, where appropriate, or back. Caches and services are named, and (Coherence) clusters are represented by their named monitoring connection.</p>
<b>OC-CACHE</b>	OcCacheTotals	Rate Store Reads	<p><b>RateStoreReads / OcCacheRateStoreReadsHigh</b></p> <p>The level of this Key Metric is <b>1</b>. (Level <b>0</b> KMs are always displayed. Level <b>1</b> KMs are displayed if the <b>Include Detail Level Metrics</b> checkbox is checked.)</p>

<b>OC-CACHE</b>	OcCacheTotals	Rate Store Writes	<p><b>RateStoreWrites / OcCacheRateStoreWritesHigh</b></p> <p>The level of this Key Metric is <b>1</b>. (Level <b>0</b> KMs are always displayed. Level <b>1</b> KMs are displayed if the <b>Include Detail Level Metrics</b> checkbox is checked.)</p>
<b>OC-CACHE</b>	OcCacheTotals	CacheSize	<p><b>CacheSize / OCCacheSizeHigh</b></p> <p>The level of this Key Metric is <b>0</b>. (Level <b>0</b> KMs are always displayed. Level <b>1</b> KMs are displayed if the <b>Include Detail Level Metrics</b> checkbox is checked.)</p> <p>This metric is the number of objects in a cache for a given tier of a given cache for a given service in a given (Coherence) cluster. The tier can be front, where appropriate, or back. Caches and services are named, and (Coherence) clusters are represented by their named monitoring connection.</p>
<b>OC-CLUSTER</b>	OcPacketStats	Packet Loss	<p><b>SentFailureRate / OcBadCommunicationCluster</b></p> <p>This metric is the (network/packet) sent failure rate averaged across all of the nodes of a cluster.</p>
<b>OC-CLUSTER NODES</b>	OcNodeTotals	CPU Used %	<p><b>AvgCpuPercent / OcClusterNodesCPUHigh</b></p> <p>This metric is the average CPU usage of all the nodes of a given storage class in a cluster. The storage class is represented by the <b>StorageEnabled</b> index column, which can be <b>true</b> or <b>false</b>. Thus metrics for storage enabled nodes in a cluster are aggregated into a cache row where <b>StorageEnabled = true</b>, and non storage enabled nodes in a cluster are aggregated into a cache row where <b>StorageEnabled = false</b>. This metric is shown as a trace in the <b>Cluster - Memory/Network Health</b> display. The metric is labeled Avg. CPU% and is displayed (for storage enabled nodes) in the Storage Nodes trend grouping and (for non storage enabled nodes) in the Process Nodes trend grouping.</p>
<b>OC-CLUSTER NODES</b>	OcNodeTotals	Packet Rx Loss	<p><b>RcvdFailureRate100 / OcClusterNodesRcvdFailureRateHigh</b></p> <p>This metric is the (network/packet) received failure rate averaged across all of the nodes of a given storage class in a cluster. The storage class is the <b>StorageEnabled</b> index column, which can be <b>true</b> or <b>false</b>. Metrics for storage enabled nodes in a cluster are aggregated into a cache row where <b>StorageEnabled = true</b>, and non storage enabled nodes in a cluster are aggregated into a cache row where <b>StorageEnabled = false</b>.</p>

<b>OC-CLUSTER NODES</b>	OcNodeTotals	Memory Used %	<b>MemoryUsedPct100 / OcClusterNodesMemHigh</b> This metric is the memory used percentage averaged across all of the nodes of a given storage class in a cluster. The storage class is the <b>StorageEnabled</b> index column, which can be <b>true</b> or <b>false</b> . Metrics for storage enabled nodes in a cluster are aggregated into a cache row where <b>StorageEnabled = true</b> , and non storage enabled nodes in a cluster are aggregated into a cache row where <b>StorageEnabled = false</b> .
<b>OC-CLUSTER NODES</b>	OcNodeTotals	Packet Tx Loss	<b>SentFailureRate100 / OcClusterNodesSentFailureRateHigh</b> This metric is the (network/packet) sent failure rate averaged across all of the nodes of a given storage class in a cluster. The storage class is the <b>StorageEnabled</b> index column, which can be <b>true</b> or <b>false</b> . Metrics for storage enabled nodes in a cluster are aggregated into a cache row where <b>StorageEnabled = true</b> , and non storage enabled nodes in a cluster are aggregated into a cache row where <b>StorageEnabled = false</b> .

### Oracle Database

The following KMs are available with the Solution Package for Oracle Database. The level of the Key Metric is **0** except where noted. Level **0** KMs are always shown in displays. Level **1** KMs are displayed if the **Include Detail Level Metrics** checkbox is selected.)

CI Type	Cache	Selector	Metric / Alert
<b>ORACLE</b>	OraDatabaseAvailability	Response Time	<b>ResponseTimeMilliSec / OraDatabaseResponseTimeHigh</b>

### Oracle WebLogic

The following KMs are available with the Solution Package for Oracle WebLogic. The level of the Key Metric is **0** except where noted. Level **0** KMs are always shown in displays. Level **1** KMs are displayed if the **Include Detail Level Metrics** checkbox is selected.

CI Type	Cache	Selector	Metric / Alert
<b>WLS</b>	WlsJvmStats	JVM CPU %	<b>JvmProcessorLoad / WlsServerCpuHigh</b>
<b>WLS</b>	WlsJvmStats	JVM Memory %	<b>MemoryUsedPercent / WlsServerMemoryUsageHigh</b>
<b>WLS</b>	WlsThreadPoolRuntime	Hogging Threads	<b>HoggingThreadCount / WlsHoggingThreadsHigh</b>
<b>WLS</b>	WlsServerRuntime	Open Sockets	<b>OpenSocketsCurrentCount / WlsOpenSocketsHigh</b> The level of this Key Metric is <b>1</b> . (Level <b>0</b> KMs are always displayed. Level <b>1</b> KMs are displayed if the <b>Include Detail Level Metrics</b> checkbox is checked.)
<b>WLS</b>	WlsThreadPoolRuntime	Thread Total Count	<b>ExecuteThreadTotalCount / WlsThreadsTotalHigh</b> The level of this Key Metric is <b>1</b> . (Level <b>0</b> KMs are always displayed. Level <b>1</b> KMs are displayed if the <b>Include Detail Level Metrics</b> checkbox is checked.)

<b>WLS-APP</b>	WlsSessionStats	Open Sessions	<b>OpenSessionsCurrentCount / WlsAppOpenSessionsHigh</b>
<b>WLS-JMS-DEST</b>	WlsJmsDestinationTotals	Messages Pending	<b>MessagesPendingCount / WlsJmsDestinationMessagesPendingHigh</b>
<b>WLS-JMS-SERVER</b>	WlsJmsServerRuntime	Messages Pending	<b>MessagesPendingCount / WlsJmsMessagesPendingHigh</b>

### RTViewManager

The following KMs are available with the RTView Manager Solution Package. The level of the Key Metric is **0** except where noted. Level **0** KMs are always shown in displays. Level **1** KMs are displayed if the **Include Detail Level Metrics** checkbox is selected.

<b>CI Type</b>	<b>Cache</b>	<b>Selector</b>	<b>Metric / Alert</b>
<b>JVM</b>	JvmOperatingSystem	Cpu %	<b>CpuPercent / JvmCpuPercentHigh</b>
<b>JVM</b>	JvmMemory	Memory %	<b>MemoryUsedPercent / JvmMemoryUsedHigh</b>
<b>JVM</b>	JvmThreading	Thread Count	<b>ThreadCount / JvmThreadCountHigh</b> The level of this Key Metric is <b>1</b> . (Level <b>0</b> KMs are always displayed. Level <b>1</b> KMs are displayed if the <b>Include Detail Level Metrics</b> checkbox is checked.)
<b>TOMCAT</b>	TomcatWebModuleTotals	Active Sessions	<b>activeSessions / TomcatActiveSessionsHigh</b>
<b>TOMCAT</b>	TomcatWebModuleTotals	Accesses / sec	<b>RateaccessCount / TomcatAccessRateHigh</b>
<b>TOMCAT-APP</b>	TomcatWebModuleStats	Active Sessions	<b>activeSessions / TomcatAppActiveSessionsHigh</b>
<b>TOMCAT-APP</b>	TomcatWebModuleStats	Accesses / sec	<b>RateaccessCount / TomcatAppAccessRateHigh</b>

### RTVRULES

The following KMs are available with the RTVRULES Solution Package which comes with RTView EM. The level of the Key Metric is **0** except where noted. Level **0** KMs are always shown in displays. Level **1** KMs are displayed if the **Include Detail Level Metrics** checkbox is selected.

<b>CI Type</b>	<b>Cache</b>	<b>Selector</b>	<b>Metric / Alert</b>
<b>EM-SERVICE</b>	RtvCmdServiceStats_local	Alert Impact	<b>AlertImpact / RtvEmServiceAlertImpactHigh</b>

## Solace

The following KMs are available with the Solution Package for Solace. The level of the Key Metric is **0** except where noted. Level **0** KMs are always shown in displays. Level **1** KMs are displayed if the **Include Detail Level Metrics** checkbox is selected.

CI Type	Cache	Selector	Metric / Alert
<b>SOLACE-MSGROUTER</b>	SolAppliances	# Msgs Spooled	<b>num-messages-spoiled / SolMsgRouterPendingMsgsHigh</b>
<b>SOLACE-MSGROUTER</b>	SolAppliances	OUT Msgs/sec	<b>total-cl-msgs-sent-per-sec / SolMsgRouterOutboundMsgRateHigh</b>
<b>SOLACE-MSGROUTER</b>	SolAppliances	IN Msgs/sec	<b>total-cl-msgs-rcvd-per-sec / SolMsgRouterInboundMsgRateHigh</b>
<b>SOLACE-VPN</b>	SolVpns	Connections	<b>connections / SolVpnConnectionCountHigh</b>

## TIBCO ActiveMatrix

The following KMs are available with the Solution Package for TIBCO ActiveMatrix. The level of the Key Metric is **0** except where noted. Level **0** KMs are always shown in displays. Level **1** KMs are displayed if the **Include Detail Level Metrics** checkbox is selected.

CI Type	Cache	Selector	Metric / Alert
<b>AMX-SERVICE</b>	AmxServiceTotals	Service Hits/Min	<b>Hits Per Minute / AmxServiceHitRateHigh</b>
<b>AMX-SERVICE</b>	AmxServiceTotals	Service Response Time	<b>Avg. Response Time / AmxServiceResponseTimeHigh</b>
<b>AMX-SERVICE NODE</b>	AmxServices	Node Hits/Min	<b>Hits Per Minute / AmxServiceNodeHitRateHigh</b>
<b>AMX-SERVICE NODE</b>	AmxServices	Node Response Time	<b>Avg. Response Time / AmxServiceNodeResponseTimeHigh</b>

## TIBCO ActiveSpaces

The following KMs are available with the Solution Package for TIBCO ActiveSpaces. The level of the Key Metric is **0** except where noted. Level **0** KMs are always shown in displays. Level **1** KMs are displayed if the **Include Detail Level Metrics** checkbox is selected.

CI Type	Cache	Selector	Metric / Alert
<b>TAS-MEMBER BYSPACE</b>	TasSeeders	Space Util by Seeder	<b>spaceUtilPerSeeder / TasMemberSeederCapacity</b>
<b>TAS-SPACE</b>	TasSpaceStatistics	Gets/sec	<b>RateGets / TasSpaceGetRateHigh</b>
<b>TAS-SPACE</b>	TasSpaceStatistics	Puts/sec	<b>RatePuts / TasSpacePutRateHigh</b>

### TIBCO BusinessEvents

The following KMs are available with the Solution Package for TIBCO BusinessEvents. The level of the Key Metric is **0** except where noted. Level **0** KMs are always shown in displays. Level **1** KMs are displayed if the **Include Detail Level Metrics** checkbox is selected.

CI Type	Cache	Selector	Metric / Alert
<b>TBE-CLUSTER</b>	TbeClusterSummary	Received Events Rate	<b>Received Events Rate / TbeClusterEventsRecvdRateHigh</b>
<b>TBE-CLUSTER</b>	TbeClusterSummary	Rules Fired Rate	<b>totalRateTotalNumberRulesFired / TbeClusterRuleFiringRateHigh</b>
<b>TBE-CLUSTER</b>	TbeClusterSummary	Concept Cache Ops Rate	<b>totalConceptOperationRate / TbeClusterConceptOpRateHigh</b>
<b>TBE-CLUSTER</b>	TbeClusterSummary	Backing Store Ops Rate	<b>totalBkngStoreOpsPerSec / TbeClusterBkngStoreOpRateHigh</b>

### TIBCO BusinessWorks (Version 5) Monitor

The following KMs are available with the Solution Package for TIBCO BusinessWorks version 5. The level of the Key Metric is **0** except where noted. Level **0** KMs are always shown in displays. Level **1** KMs are displayed if the **Include Detail Level Metrics** checkbox is selected.

CI Type	Cache	Selector	Metric / Alert
<b>BW-ENGINE</b>	BwEngines	CPU Used %	<b>CPU % / BwEngineCpuUsedHigh</b>
<b>BW-ENGINE</b>	BwEngines	Memory Used %	<b>PercentUsed / BwEngineMemUsedHigh</b>
<b>BW-PROCESS</b>	BwProcesses	AverageElapsed	<b>Process Avg Elapsed Time / BwProcessAvgElapsedTimeHigh</b>
<b>BW-PROCESS</b>	BwProcesses	RateCreated / sec	<b>Processes Created/sec / BwProcessCreatedRateHigh</b>
<b>BW-PROCESS</b>	BwProcesses	TotalCpuPercent	<b>Process Total CPU Percent / BwProcessTotalCpuPercentHigh</b>
<b>BW-PROCESS</b>	BwProcesses	Process Exec Time / sec	<b>RateTotalExecution / BwProcessExecutionTimeHigh</b>
<b>BW-SERVER</b>	BwServers	CPU Used %	<b>CPU Usage % / BwServerCpuUsedHigh</b>

### TIBCO BusinessWorks (Version 6) Monitor

The following KMs are available with the Solution Package for TIBCO BusinessWorks version 6. The level of the Key Metric is **0** except where noted. Level **0** KMs are always shown in displays. Level **1** KMs are displayed if the **Include Detail Level Metrics** checkbox is selected.

CI Type	Cache	Selector	Metric / Alert
<b>BW6-APPNODE</b>	Bw6AppNodes	CPU Used %	<b>Used CPU Percentage / Bw6AppNodeCpuUsedHigh</b>
<b>BW6-APPNODE</b>	Bw6AppNodes	Mem Used %	<b>Used Memory Percentage / Bw6AppNodeMemUsedHigh</b>
<b>BW6-APP</b>	Bw6ProcessTotalsByApp	App Created / sec	<b>RateCreated / Bw6AppProcessCreatedRateHigh</b>
<b>BW6-APP</b>	Bw6ProcessTotalsByApp	App Exec Time / sec	<b>RateTotal Execution / Bw6AppProcessExecutionTimeHigh</b>
<b>BW6-PROCESS</b>	Bw6Processes	Process Created / sec	<b>RateCreated / Bw6ProcessCreatedRateHigh</b>
<b>BW6-PROCESS</b>	Bw6Processes	Process Exec Time / sec	<b>RateTotal Execution / Bw6ProcessExecutionTimeHigh</b>

### TIBCO EMS

The following KMs are available with the Solution Package for TIBCO EMS. The level of the Key Metric is **0** except where noted. Level **0** KMs are always shown in displays. Level **1** KMs are displayed if the **Include Detail Level Metrics** checkbox is selected.

CI Type	Cache	Selector	Metric / Alert
<b>EMS-QUEUE</b>	EmsQueues	Pending Msgs	<b>pendingMessageCount / EmsQueuesPendingMsgsHigh</b>
<b>EMS-QUEUE</b>	EmsQueues	In Msgs / sec	<b>inboundMessageRate / EmsQueuesInMsgRateHigh</b>
<b>EMS-QUEUE</b>	EmsQueues	Out Msgs / sec	<b>outboundMessageRate / EmsQueuesOutMsgRateHigh</b>
<b>EMS-QUEUE</b>	EmsQueues	Consumers	<b>consumerCount / EmsQueuesConsumerCountHigh</b> The level of this Key Metric is <b>1</b> . (Level <b>0</b> KMs are always displayed. Level <b>1</b> KMs are displayed if the <b>Include Detail Level Metrics</b> checkbox is checked.)
<b>EMS-SERVER</b>	EmsServerInfo	Pending Msgs	<b>pendingMessageCount / EmsServerPendingMsgsHigh</b>
<b>EMS-SERVER</b>	EmsServerInfo	In Msgs / sec	<b>inboundMessageRate / EmsServerInMsgRateHigh</b>
<b>EMS-SERVER</b>	EmsServerInfo	Out Msgs / sec	<b>outboundMessageRate / EmsServerOutMsgRateHigh</b>

<b>EMS-SERVER</b>	EmsServerInfo	Msg Mem %	<b>messageMemoryPct / EmsServerMemUsedHigh</b> The level of this Key Metric is <b>1</b> . (Level <b>0</b> KMs are always displayed. Level <b>1</b> KMs are displayed if the <b>Include Detail Level Metrics</b> checkbox is checked.)
<b>EMS-SERVER</b>	EmsServerInfo	Connections	<b>connectionCount / EmsServerConnectionCountHigh</b> The level of this Key Metric is <b>1</b> . (Level <b>0</b> KMs are always displayed. Level <b>1</b> KMs are displayed if the <b>Include Detail Level Metrics</b> checkbox is checked.)
<b>EMS-SERVER</b>	EmsServerInfo	Async DB Size	<b>asyncDBSize / EmsServerAsyncDBSizeHigh</b> The level of this Key Metric is <b>1</b> . (Level <b>0</b> KMs are always displayed. Level <b>1</b> KMs are displayed if the <b>Include Detail Level Metrics</b> checkbox is checked.)
<b>EMS-SERVER</b>	EmsServerInfo	Sync DB Size	<b>syncDBSize / EmsServerSyncDBSizeHigh</b> The level of this Key Metric is <b>1</b> . (Level <b>0</b> KMs are always displayed. Level <b>1</b> KMs are displayed if the <b>Include Detail Level Metrics</b> checkbox is checked.)
<b>EMS-TOPIC</b>	EmsTopics	Pending Msgs	<b>pendingMessageCount / EmsTopicsPendingMsgsHigh</b>
<b>EMS-TOPIC</b>	EmsTopics	In Msgs / sec	<b>inboundMessageRate / EmsTopicsInMsgRateHigh</b>
<b>EMS-TOPIC</b>	EmsTopics	Out Msgs / sec	<b>outboundMessageRate / EmsTopicsOutMsgRateHigh</b>
<b>EMS-TOPIC</b>	EmsTopics	Consumers	<b>consumerCount / EmsTopicsConsumerCountHigh</b> The level of this Key Metric is <b>1</b> . (Level <b>0</b> KMs are always displayed. Level <b>1</b> KMs are displayed if the <b>Include Detail Level Metrics</b> checkbox is checked.)
<b>EMS-TOPIC</b>	EmsTopics	Subscribers	<b>subscriberCount / EmsTopicsSubscriberCountHigh</b> The level of this Key Metric is <b>1</b> . (Level <b>0</b> KMs are always displayed. Level <b>1</b> KMs are displayed if the <b>Include Detail Level Metrics</b> checkbox is checked.)

## UX

The following KMs are available with the Solution Package for UX. The level of the Key Metric is **0** except where noted. Level **0** KMs are always shown in displays. Level **1** KMs are displayed if the **Include Detail Level Metrics** checkbox is selected.

<b>CI Type</b>	<b>Cache</b>	<b>Selector</b>	<b>Metric / Alert</b>
<b>UX-URL</b>	UXURLData	Response Time	<b>MostRecentTime / UXURLResponseSlow</b>

## VMWare vSphere

The following KMs are available with the Solution Package for VMWare vSphere. The level of the Key Metric is **0** except where noted. Level **0** KMs are always shown in displays. Level **1** KMs are displayed if the **Include Detail Level Metrics** checkbox is selected.

CI Type	Cache	Selector	Metric / Alert
<b>VMWARE -HOST</b>	VmwHostSystems	CPU Usage	<b>cpu.usage.average / VmwHostCpuUtilizationHigh</b>
<b>VMWARE -HOST</b>	VmwHostSystems	Memory Usage	<b>mem.usage.average / VmwHostMemoryUsageHigh</b>
<b>VMWARE -VM</b>	VmwVirtualMachines	CPU Usage	<b>cpu.usage.average / VmwVmCpuUtilizationHigh</b>
<b>VMWARE -VM</b>	VmwVirtualMachines	Memory Usage	<b>mem.usage.average / VmwVmMemoryUsageHigh</b>

## Component Views

These displays present the lowest level view of CMDB contents--the component level. In these displays, alert states for components are shown by Service and Area in tabular and heatmap formats, while highlighting the most critical alert state for each component. Data can be filtered by Areas, Services, Groups, Regions and Environment. Data is filtered by the \$rtvOwnerMask, \$rtvAreaMask, \$rtvGroupMask and \$rtvServiceMask values for the logged in user. For details, refer to the *RTView Enterprise Configuration Guide*.

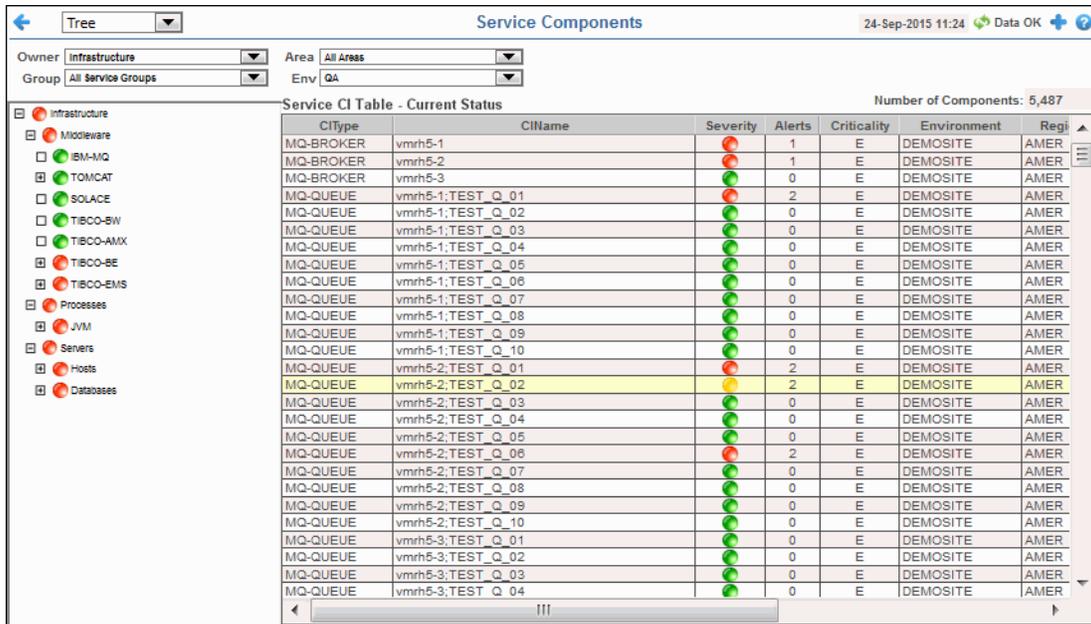
Use these displays to determine whether a component is malfunctioning. Displays in this View are:

- "[CI / Service Tree View](#)": Table of CMDB contents for all component-level details by Service for all Owners, Areas, Groups, Regions and Environments (without the option to filter).
- "[CI / Service Table](#)": Table of CMDB contents for all component-level details by Service for all Owners, Areas, Groups, Regions and Environments (without the option to filter).
- "[CI / Type Heatmap](#)": Heatmap of CMDB contents organized by CIType, with the option to filter by Owner, Area, Group, Environment and alert Metric, and show CI Names.
- "[CI / Type Table](#)": Table of CMDB contents for all component-level details for all Areas, Services, Groups, Regions and Environments, with the option to filter by Owner and one or all Areas, Groups and Environments.

### CI / Service Tree View

View the contents of the CMDB hierarchically ordered in a navigation tree. Each row in the table is a different CI (for example, **localhost;RTVMGR\_DATASERVER**).

Make a selection from the **Owner** drop-down menu, then use the navigation tree to filter data in the **Service CI Table**. The navigation tree, which provides a visual of the CMDB hierarchy, provides further filtering to the **Area**, **Group**, and **Environment** drop-down menus. Click Sort  to order column data.



Service Components 24-Sep-2015 11:24 Data OK

Owner: Infrastructure Area: All Areas  
 Group: All Service Groups Env: QA

Service CI Table - Current Status Number of Components: 5,487

CIType	CIName	Severity	Alerts	Criticality	Environment	Regi
MQ-BROKER	vmrh5-1		1	E	DEMOSITE	AMER
MQ-BROKER	vmrh5-2		1	E	DEMOSITE	AMER
MQ-BROKER	vmrh5-3		0	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-1;TEST_Q_01		2	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-1;TEST_Q_02		0	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-1;TEST_Q_03		0	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-1;TEST_Q_04		0	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-1;TEST_Q_05		0	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-1;TEST_Q_06		0	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-1;TEST_Q_07		0	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-1;TEST_Q_08		0	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-1;TEST_Q_09		0	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-1;TEST_Q_10		0	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-2;TEST_Q_01		2	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-2;TEST_Q_02		2	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-2;TEST_Q_03		0	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-2;TEST_Q_04		0	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-2;TEST_Q_05		0	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-2;TEST_Q_06		2	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-2;TEST_Q_07		0	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-2;TEST_Q_08		0	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-2;TEST_Q_09		0	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-2;TEST_Q_10		0	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-3;TEST_Q_01		0	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-3;TEST_Q_02		0	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-3;TEST_Q_03		0	E	DEMOSITE	AMER
MQ-QUEUE	vmrh5-3;TEST_Q_04		0	E	DEMOSITE	AMER

**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu , Table open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

**Row Color Code:**

Tables with colored rows indicate the following:

- Red indicates that one or more alerts exceeded their ALARM LEVEL threshold in the table row.
- Yellow indicates that one or more alerts exceeded their WARNING LEVEL threshold in the table row.
- Green indicates that no alerts exceeded their WARNING or ALARM LEVEL threshold in the table row.

**Filter By:**

The following filtering options are typically included:

**Owner:** Choose an Owner to see metrics for Areas associated with that Owner.

**Area:** Choose an Area to see metrics for Groups associated with that Area and Owner.

**Group:** Choose a Group to see metrics for Services associated with that Group, Area and Owner.

**Service:** Choose a Service to see metrics for Environments associated with that Service, Group, Area and Owner.

**Env:** Choose an Environment to see metrics for Environments associated with that Service, Group, Area and Owner.

**Fields and Data**

This display includes:

<b>Number of Components</b>	The total number of CIs currently in the table.
<b>CIType</b>	The type of CI.
<b>CIName</b>	The name or address of the CI.
<b>Severity</b>	The maximum level of alerts for the CI. Values range from 0 to 2, where 2 is the greatest Alert Severity: <ul style="list-style-type: none"> <li><span style="color: red;">●</span> One or more alerts exceeded their ALARM LEVEL threshold.</li> <li><span style="color: yellow;">●</span> One or more alerts exceeded their WARNING LEVEL threshold.</li> <li><span style="color: green;">●</span> No alert thresholds have been exceeded.</li> </ul>
<b>Criticality</b>	The Criticality (rank of importance) specified in the Service Data Model (CMDB) by your administrator. Criticality values are listed in the Component Views / CI Service Table display, which range from A to E, where A is the highest Criticality. This value is used to determine the value for Alert Impact.
<b>Environment</b>	The Environment for the CI.
<b>Region</b>	The name of the Region for the CI.
<b>City</b>	The name of the City for the CI.
<b>Country</b>	The name of the Country for the CI.
<b>SiteName</b>	The name of the Site for the CI.
<b>OSType</b>	The operating system currently running on the CI.
<b>City</b>	The name of the City for the CI.
<b>Country</b>	The name of the Country for the CI.

**CI / Service Table**

View the contents of the CMDB, without filtering, in a tabular format. Each row in the table is a different CI (for example, **localhost;RTVMGR\_DATASERVER**).

Use the available drop-down menus or right-click to filter data shown in the display. Click Sort  to order column data.



Owner	Area	ServiceGroup	ServiceName	CType	CIName	Se
Jerelyn Parker	Backends	IBM	MQ	MQ-BROKER	vmrh5-1	
Jerelyn Parker	Backends	IBM	MQ	MQ-BROKER	vmrh5-2	
Jerelyn Parker	Backends	IBM	MQ	MQ-BROKER	vmrh5-3	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-1;TEST_Q_01	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-1;TEST_Q_02	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-1;TEST_Q_03	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-1;TEST_Q_04	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-1;TEST_Q_05	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-1;TEST_Q_06	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-1;TEST_Q_07	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-1;TEST_Q_08	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-1;TEST_Q_09	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-1;TEST_Q_10	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-2;TEST_Q_01	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-2;TEST_Q_02	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-2;TEST_Q_03	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-2;TEST_Q_04	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-2;TEST_Q_05	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-2;TEST_Q_06	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-2;TEST_Q_07	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-2;TEST_Q_08	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-2;TEST_Q_09	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-2;TEST_Q_10	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-3;TEST_Q_01	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-3;TEST_Q_02	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-3;TEST_Q_03	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-3;TEST_Q_04	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-3;TEST_Q_05	
Jerelyn Parker	Backends	IBM	MQ	MQ-QUEUE	vmrh5-3;TEST_Q_06	

**Title Bar (possible features are):**

-   Open the previous and upper display.
-  Open an instance of this display in a new window.
-  Open the online help page for this display.
- open commonly accessed displays.

 **Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Row Color Code:**

Tables with colored rows indicate the following:

-  Red indicates that one or more alerts exceeded their ALARM LEVEL threshold in the table row.
-  Yellow indicates that one or more alerts exceeded their WARNING LEVEL threshold in the table row.
-  Green indicates that no alerts exceeded their WARNING or ALARM LEVEL threshold in the table row.

**Filter By:**

The following filtering options are typically included:

- Owner:** Choose an Owner to see metrics for Areas associated with that Owner.
- Area:** Choose an Area to see metrics for Groups associated with that Area and Owner.
- Group:** Choose a Group to see metrics for Services associated with that Group, Area and Owner.
- Service:** Choose a Service to see metrics for Environments associated with that Service, Group, Area and Owner.
- Env:** Choose an Environment to see metrics for Environments associated with that Service, Group, Area and Owner.

**Fields and Data**

This display includes:

<b>Number of Rows</b>	The current total number of rows in the table.
<b>Service CI Table</b>	
<b>Owner</b>	The Owner the CI is associated with.
<b>Area</b>	The Area the CI is associated with.
<b>ServiceGroup</b>	The Group the CI is associated with.
<b>ServiceName</b>	The Service the CI is associated with.
<b>CIType</b>	The type of CI.
<b>CIName</b>	The name or address of the CI.
<b>Severity</b>	<p>The maximum level of alerts for the CI. Values range from <b>0</b> to <b>2</b>, where <b>2</b> is the greatest Alert Severity:</p> <ul style="list-style-type: none"> <li><span style="color: red;">●</span> One or more alerts exceeded their ALARM LEVEL threshold.</li> <li><span style="color: yellow;">●</span> One or more alerts exceeded their WARNING LEVEL threshold.</li> <li><span style="color: green;">●</span> No alert thresholds have been exceeded.</li> </ul>
<b>Criticality</b>	The Criticality (rank of importance) specified in the Service Data Model (CMDB) by your administrator. Criticality values are listed in the <b>Component Views - CI Service Table</b> display, which range from <b>A</b> to <b>E</b> , where <b>A</b> is the highest Criticality. This value is used to determine the value for Alert Impact.
<b>Environment</b>	The Environment for the CI.
<b>City</b>	The name of the City for the CI.
<b>Country</b>	The name of the Country for the CI.
<b>Region</b>	The name of the Region for the CI.
<b>SiteName</b>	The name of the Site for the CI.

**CI / Type Heatmap**

View heatmap of alert states for CIs in all or one Area, Group or Environment. The heatmap organizes CIs by CI Type, and uses color to show the most critical alert state for each. Each rectangle in the heatmap represents a CI (for example, **localhost;RTVMGR\_DATASERVER**).

Use the available drop-down menus or right-click to filter data shown in the display. Use the check-boxes  to include or exclude labels in the heatmap. Move your mouse over a rectangle to see additional information. Double-click (or right-click and select **Drill Down**) a rectangle in the heatmap to view details relevant to the CI Type. By default, this display shows all Areas, Groups, and Environments and alert Impact.



#### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

#### Filter By:

The following filtering options are typically included:

**Owner:** Choose an Owner to see metrics for Areas associated with that Owner.

**Area:** Choose an Area to see metrics for Groups associated with that Area and Owner.

**Group:** Choose a Group to see metrics for Services associated with that Group, Area and Owner.

**Service:** Choose a Service to see metrics for Environments associated with that Service, Group, Area and Owner.

**Env:** Choose an Environment to see metrics for Environments associated with that Service, Group, Area and Owner.

#### Metric:

Choose the type of metric to show in the heatmap. Each metric has its own gradient bar that maps relative values to colors:

##### Alert Impact

The product of the maximum Alert Severity of alerts in the heatmap rectangle multiplied by the maximum Criticality of alerts in the heatmap rectangle. Values range from **0** - **10**, as indicated in the color gradient bar, where **10** is the highest Alert Impact.

**Alert Severity** The maximum level of alerts in the heatmap rectangle. Values range from **0** - **2**, as indicated in the color gradient  bar, where **2** is the highest Alert Severity.

-  Red indicates that one or more metrics have reached their alarm threshold. Metrics that have exceeded their specified ALARM LEVEL threshold have an Alert Severity value of **2**.
-  Yellow indicates that one or more metrics have reached their alarm threshold. Metrics that have exceeded their specified WARNING LEVEL threshold have an Alert Severity value of **1**.
-  Green indicates that no metrics have reached their alert thresholds. Metrics that have not exceeded their specified thresholds have an Alert Severity value of **0**.

**Alert Count** The total number of critical and warning alerts in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.

**Criticality** The maximum level of Criticality (rank of importance) in the heatmap rectangle. Values range from **1** to **5**, as indicated in the color gradient  bar, where **5** is the highest Criticality.

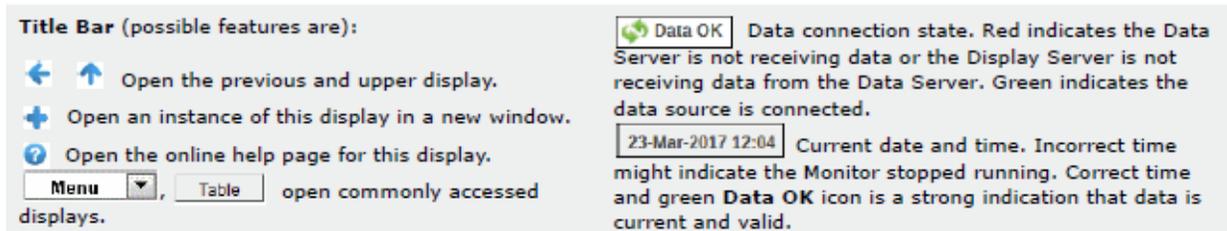
Criticality is specified in the Service Data Model (CMDB) by your administrator. Criticality values are listed in the **Component Views - "CI / Service Table"** display, which range from **A** to **E**, where **A** is the highest Criticality (level **5** maps to a Criticality of **A** and level **1** maps to a Criticality of **E** with equally spaced intermediate values).

### CI / Type Table

View tabular list of all CIs by CType, as well as their alert metrics (Impact, Severity and Count, for one or all Areas, Groups or Environments). Each row in the table is a different CI (for example, **localhost;RTVMGR\_DATASERVER**). The row color represents the most critical alert state for the CI.

Use the available drop-down menus or right-click to filter data shown in the display. Click Sort  to order column data.

Table		All Components By Type			24-Sep-2015 11:16	Data OK	
Owner:	Infrastructure	Area:	All Areas	CIs:	64	Env:	QA
Group:	All Service Groups						
CType	CIName	Severity	AlertCount	AlertImpact			
EMS-QUEUE	top://192.168.200.132:7222;queue.sample		3	2			
EMS-QUEUE	top://192.168.200.132:7222;sample		3	2			
EMS-SERVER	top://192.168.200.132:7222		2	2			
EMS-TOPIC	top://192.168.200.132:7222;sample		2	2			
EMS-TOPIC	top://192.168.200.132:7222;topic.sample		2	2			
EMS-TOPIC	top://192.168.200.132:7222;topic.sample.exported		2	2			
EMS-TOPIC	top://192.168.200.132:7222;topic.sample.imported		2	2			
HOST	QATB;SLHOST-WIN3		2	2			
HOST	QATB;SLHOST-WIN4		1	1			
JVM	localhost;ALERTHISTORIAN		1	1			
JVM	localhost;ALERT_SERVER		3	2			
JVM	localhost;AMXMON-HISTORIAN		1	2			
JVM	localhost;AMXMON-SLHOST-WIN3		2	1			
JVM	localhost;AMXMON-SLHOST-WIN4		2	1			
JVM	localhost;BWMON-HISTORIAN		1	2			
JVM	localhost;BWMONITOR-WIN-8		1	2			
JVM	localhost;CONFIG_SERVER		2	1			
JVM	localhost;DISPLAYSERVER		2	1			
JVM	localhost;DISPLAYSERVER_DARKSTYLES		2	1			
JVM	localhost;EMSMON-HISTORIAN		1	2			
JVM	localhost;EMSMON-SLHOST-WIN3		2	2			
JVM	localhost;EMSMON-SLHOST-WIN4		2	1			
JVM	localhost;EMSMONITOR-WIN-8		2	2			
JVM	localhost;MISCMON-HISTORIAN		1	2			

**Row Color Code:**

Tables with colored rows indicate the following:

- Red indicates that one or more alerts exceeded their ALARM LEVEL threshold in the table row.
- Yellow indicates that one or more alerts exceeded their WARNING LEVEL threshold in the table row.
- Green indicates that no alerts exceeded their WARNING or ALARM LEVEL threshold in the table row.

**Filter By:**

The following filtering options are typically included:

**Owner:** Choose an Owner to see metrics for Areas associated with that Owner.

**Area:** Choose an Area to see metrics for Groups associated with that Area and Owner.

**Group:** Choose a Group to see metrics for Services associated with that Group, Area and Owner.

**Service:** Choose a Service to see metrics for Environments associated with that Service, Group, Area and Owner.

**Env:** Choose an Environment to see metrics for Environments associated with that Service, Group, Area and Owner.

**Fields and Data**

This display includes:

**CI Count** The total number of CIs listed in the table. This value is determined by the selections made from display drop-down menus. The totals number for each Environment are also shown.

**CI Table**

This table lists all CIs for the selected Group. Each row in the table is a CI. Each CI can have multiple alerts. Click a CI to view alerts for the CI in the lower table.

**CIType** The type of CI.

**CIName** The name or address of the CI.

**Severity** The maximum level of alerts for the CI. Values range from 0 to 2, where 2 is the greatest Alert Severity:

- One or more alerts exceeded their ALARM LEVEL threshold.
- One or more alerts exceeded their WARNING LEVEL threshold.
- No alert thresholds have been exceeded.

**Alert Count** The total number of critical and warning alerts for the CI.

**Alert Impact** The product of the maximum Alert Severity multiplied by the maximum Criticality of alerts. Values range from **0** - **10**, where **10** is the highest Alert Impact.

## Metric Explorer

The Metric Explorer (MX) is a tool for creating and viewing custom dashboards, referred to as *MX Views*. An MX View contains a trend graph with up to five traces which you can configure to show numeric metrics from any EM solution package. While EM provides out-of-the-box Views of metric data, there might not be a single display that shows all the metrics that are critical to a single application. MX allows end-users to create Views containing the metrics that are important to them. The MX Views your end-users create are accessed from the MX **View** drop-down menu (rather than the navigation tree as RTView Enterprise Views are). Data is filtered by the \$rtvOwnerMask, \$rtvAreaMask, \$rtvGroupMask and \$rtvServiceMask values for the logged in user. For details, refer to the *RTView Enterprise Configuration Guide*.

Displays in this View are:

- “Metric Explorer”:

---

**Note:** The Metric Explorer was added in RTView Enterprise version 1.5.0. For instructions about adding the Metric Explorer to applications created with versions older than 1.5.0, see the RTView Enterprise Upgrade Notes.

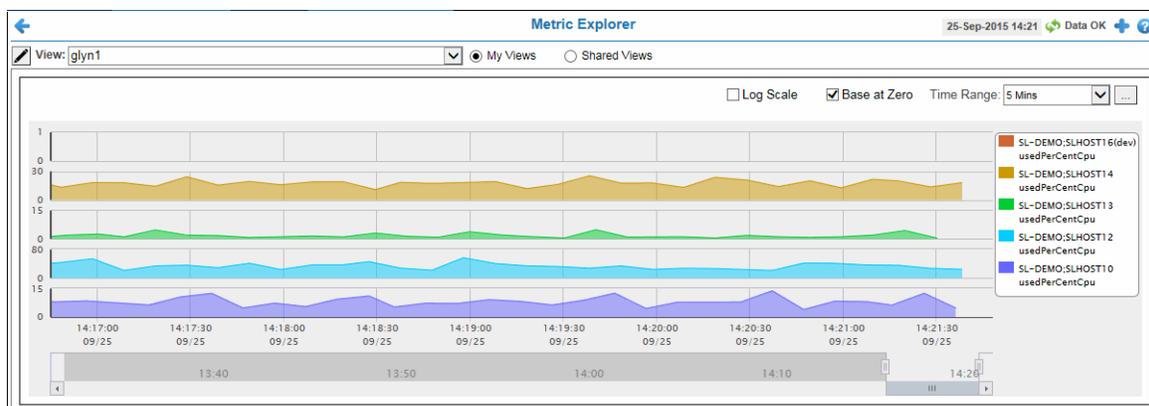
---

## Metric Explorer

View your previously created MX Views. Select an MX View from the **View** drop-down menu. The contents of the **View** drop-down menu depend on whether you choose **My Views** or **Shared Views**. Choose **My Views** to see public and private MX Views owned by you. Choose **Shared Views** to see public MX Views owned by you and other users. A public MX View is an MX View where the creator chose the **Share View with Others** option. The creator of the MX View is the owner.

Each MX View has options to apply **Log Scale**, **Base at Zero** and **Time Range** to your graphs.

To create or edit an MX View click Edit  to open the edit pane. For details, see “[Creating MX Views](#)”.



**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** , **Table** open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

## Fields and Data

Options include:



Open the edit pane.

### View

Select an MX View from the **View** drop-down menu.

### My Views

Choose **My Views** to see public and private MX Views owned by you in the **View** drop-down menu.

### Shared View

Choose **Shared Views** to see public MX Views owned by you and other users. A public MX View is an MX View where the creator chose the **Share View with Others** option. The creator of the MX View is the owner.

### Log Scale

Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.

### Base at Zero

Use zero as the Y axis minimum for all graph traces.

### Time Range

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .

The dialog box titled "Select or Enter Date and Time:" contains a text input field with a calendar icon on the right. Below the input field are two blue navigation arrows (left and right). Underneath the arrows is a button labeled "Restore to Now". At the bottom of the dialog are three buttons: "Ok", "Apply", and "Cancel".

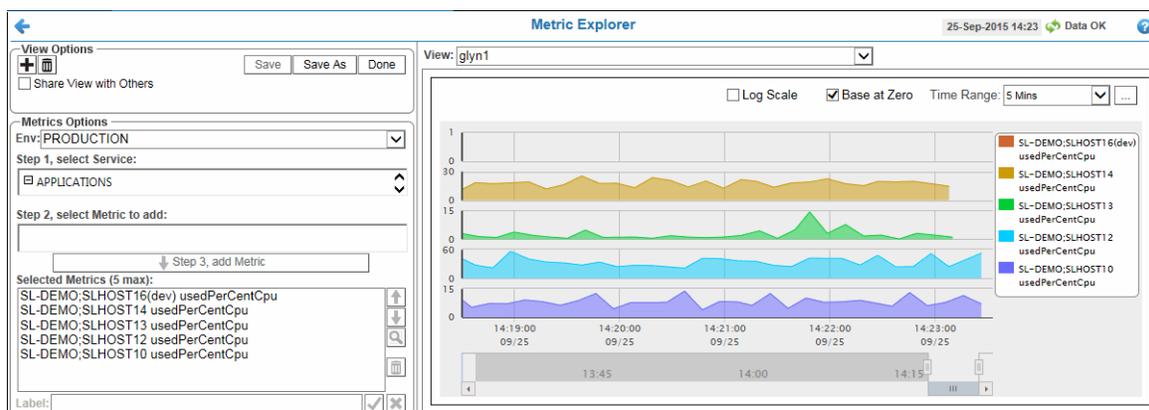
By default, the time range end point is the current time. To change the time range end point, click Calendar and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Creating MX Views

Click Edit  to open the edit pane. If an MX View is already selected, click New  to start a new MX View.



Select the Service containing the metric you want to see from the **Metrics Options/Service Tree**. The **Service Tree** is filtered by the environment in the **Env** menu. If necessary select the environment containing your Service. When you select a Service, all available metrics for that Service are listed in the **Metric Tree**. Because the metric will be displayed in a trend graph, only numeric metrics with history are listed in the **Metric Tree**. Select the metric you want to see from the **Metrics Options/Metric Tree** and click **Add Metric**. The metric is added to the **Selected Metrics** list and the MX View preview (in the right panel). Add up to five metrics to your MX View.

To change the order in which the metrics are displayed in the graph use the Up  and Down  arrows. To remove a metric, select it in the **Selected Metrics** and click Trash . To add a label to your metric, select it in the **Selected Metrics** list and enter your label text in the **Label** field. Click Apply  to apply the label, or Cancel  to cancel the label.

Click **Save** and enter a descriptive MX View name. Click **Share View with Others** to make your MX View public, otherwise, the MX View is only available to you. Click Confirm  to write the MX View to the database. Click Cancel  to return to edit mode. Click **Done** to return to the **Metric Explorer** page. The MX View you created is added to the **View** drop-down menu.

To create a new MX View with the Service already selected, select a Service from a **Service Summary Views** display and click MX  (or the table context menu). This opens the MX edit pane with the Service already selected in the MX edit pane **Service Tree**. If you selected a CI Type or CI, these are also already selected in the MX edit pane **Metric Tree**. This spares you from having to search for the Service, CI Type or CI in the **Service** and **Metrics Trees**. The displays from which you can use this feature are:

- **Service Summary Views** - "Service By CI Type"
- **Service Summary Views** - "Service Summary"

## Editing MX Views

In the Metric Explorer, select the MX View you want to edit and click Edit . The edit pane opens with the selected MX View in edit mode. To delete the MX View click Trash . To save your MX View under a new name, click **Save As**. Add, remove, reorder or label metrics as described in the "**Creating MX Views**" section (above). Select a metric in the **Selected Metrics** list and click on **Search** to update the selection in the **Service Tree** and **Metric Tree** to the values used when that metric was added to the MX View. This is useful when you want to see which Service contains a metric so you can add more metrics from the same Service.

When you are finished editing your metric, you can click **Cancel** to cancel your changes or **Save** to save your changes. To edit another MX View, select it from the **View** drop-down menu. Click **Done** to return to the Metric Explorer page.

---

**Note:** When you edit an MX View you do not own a copy of the MX View is automatically created and you are prompted to enter a name for the MX View when you save it.

---

### View Options

	Create a new MX View.
	Delete the selected MX View.
<b>Save</b>	Save the selected MX View. If this is an existing MX View, the save is done immediately. If this is a new MX View, the <b>Name</b> field becomes available and you must enter a name and click <b>Confirm Save</b> to save your MX View.
<b>Save As</b>	Save the selected MX View under a new name. The <b>Name</b> field becomes available and you must enter a name and click <b>Confirm Save</b> to save your MX View.
<b>Done</b>	Close the edit pane. This option is available when you do not have unsaved changes.
<b>Cancel</b>	Cancel your edits.
<b>Name</b>	Enter a name for your MX View. This field is available when saving a new MX View or after you click <b>Save As</b> .
<input checked="" type="checkbox"/>	Confirm that you want to save your MX View after you enter a name. This option is available when saving a new MX View or after you click <b>Save As</b> .
	Cancel the save. This is available when saving a new MX View or after you click <b>Save As</b> .
<b>Share View</b>	Select to make your MX View public. Public MX Views are available to all users in the <b>View</b> drop-down menu when the <b>Shared Views</b> option is selected. Deselect to make this MX View only available to you.

### Metric Options

<b>Env</b>	Select an Environment to filter the items in the Service Tree.
<b>Service Tree</b>	The CMDB service model (Owner, Area, Group, Service). Select a Service to populate the Metric Tree with metrics for that Service. The Services in the Service Tree are filtered by the following login substitutions: <b>\$rtvOwnerMask</b> , <b>\$rtvAreaMask</b> , <b>\$rtvGroupMask</b> and <b>\$rtvServiceMask</b> . For details, refer to the <i>RTView Enterprise Configuration Guide</i> .
<b>Metric Tree</b>	The available metrics for the selected service. The tree hierarchy is CI Type, CI name, Metric (cache: metric). The tree only contains numeric metrics with history.
<b>Add Metric</b>	Add the selected metric to the MX View. When a metric is added to the MX View, it appears in the graph.
<b>Selected Metrics</b>	The list of metrics for this MX View.

	Move the metric up in the list of selected metrics.
	Move the metric down in the list of selected metrics.
	Set the selection in the Service and Metric trees to the values used when you added the selected metric to the MX View. <b>Note:</b> If your CMDB has changed such that the Service you used to add this metric no longer exists, the search button will fail
	Delete the selected metric from the MX View.
<b>Label</b>	Enter a label to use for the selected metric. This label is not applied until you click on the confirm label button. This label is used in the graph legend.
	Confirm the label you entered for the selected metric.
	Discard the label you entered for the selected metric (revert back to the previously applied value).

## Limitations

- The Search  button fails without an error if the Service that was selected when you initially added the metric is no longer in your CMDB. To fix this, delete the metric and add it again from a Service that is currently in your CMDB. **Note:** The missing Service only makes the Search button fail. It does not cause any problems with viewing the metric.
- When you try to add a metric to an MX View that already contains that metric, it will not be added again. In the Viewer, an error message will come up saying that the metric is already in the MX View. In the Thin Client, no error is shown.
- MX Views are limited to five metrics. After a view contains five metrics, the **Add Metric** button is disabled.
- There is no indicator that shows if the MX database or Central Configuration Server are off-line in the MX configuration display. Any changes you save when either the MX database or Central Configuration Server are off line will be lost.
- When you save an MX View, RTView writes to both the View Table and the Metrics Table to the database even if only one or the other changed.
- When you save an MX View, the MX Configuration UI temporarily reverts back to the previous version of the MX View for one update, then updates with the latest changes.
- By default, MX attaches to the history\_combo table for the metric history. If the cache is not configured with a history\_combo table, the Metric Explorer will instead make a one-time attachment to the history table. In this case, toggling the **Log Scale** check-box will cause all points plotted after the initial history query to be lost. On the next update of current data a straight line will be drawn from the last history point to the new current data point.

## Alert Views

These displays present detailed information about all alerts that have occurred in your RTView Enterprise system (all Owners and all Areas). The type of alerts that appear in these displays depends on the solution packages installed on your RTView Enterprise system. Displays in this View are:

- “RTView Alerts Table”: Shows current alert data. Use this time-ordered tabular view to track, manage and assign alerts.
- “Alerts Table - HTML”: The HTML version of the RTView Alerts Table shows current alert data and is used to track, manage and assign alerts.
- “Alert History Table”: Shows historical alert data. Use this time-ordered tabular view to track alert status changes.

### RTView Alerts Table

Use this display to track and manage all alerts that have occurred in the system, add comments, acknowledge or assign Owners to alerts.

The color coded navigation tree shows the contents of the CMDB hierarchically ordered. Choose a node to filter alerts shown in the table. The **Alerts Table** only shows alerts associated with the node you select. A green indicator means the node has no associated alerts. A red indicator means the node has one or more associated alerts.

Service name labels are appended with the Environment and number of alerts. For example, the following illustrates that the **TBE** Service currently has no (**0**) associated alerts in the **PRODUCTION** Environment.

▼  TIBCO-AS  
      TAS-MEMBER (PRODUCTION)

Each row in the table is a different active alert. Select one or more rows, right-click and choose **Alert** to see all actions that you can perform on the selected alert(s). Choose **Alert / Set Filter Field** to apply the selected cell data to the **Field Filter** and **Search Text** fields. Or enter filter criteria directly in the **Field Filter** and **Search Text** fields. Click **Clear** to clear the **Field Filter** and **Search Text** fields.

Click a column heading to sort the table on that column data.

Optionally, you can use the **\$rtvUserShowDualTables** substitution to add a table that lists alerts owned by the logged in user.

The screenshot shows the 'Alerts Table' interface. At the top, there are navigation buttons (back, forward), a 'Current' dropdown, an 'Alert Group' dropdown set to 'All', and a 'Data OK' status indicator. Below this is a 'Field Filter' section with a search text box and a 'RegExOwner Filter' dropdown set to 'All'. A 'CMDB Filter' section contains a text box with the filter expression 'Owner = \* | Area = \* | Group = \* | Service = \* | Env = \*'. The main table has columns: 'Total' (166 / 166), 'Critical' (164 / 164), 'Warning' (2 / 2), and 'Suppressed' (0 / 0). The table headers are: 'First Occ', 'Last Occ', 'Count', 'Sup', 'Owner', 'Alert Name', 'Primary Service', and 'CI'. The table contains multiple rows of alert data, with some rows highlighted in red (indicating alarm level exceeded) and others in yellow (indicating warning level exceeded). At the bottom, there are 'Columns' checkboxes for 'Id', 'Closed', 'Closed Reason', and 'Alert Index', along with 'Go To CI', 'Options', and 'Details' buttons.

#### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

The row color indicates the following:

#### Row Color Code:

Tables with colored rows indicate the following:

- Red indicates that one or more alerts exceeded their ALARM LEVEL threshold in the table row.
- Yellow indicates that one or more alerts exceeded their WARNING LEVEL threshold in the table row.
- Green indicates that no alerts exceeded their WARNING or ALARM LEVEL threshold in the table row.
- Gray indicates that the alert engine that is hosting the alert is not connected, not enabled or not initialized. When you select a gray row the **Own**, **Suppress**, **Unsuppress**, **Close**, **Annotate**, **Options** and **Details** options are disabled.

**Fields and Data**

This display includes:

<b>Field Filter</b>	<p>Select a table column from the drop-down menu to perform a search in: <b>Alert Name, Alert Text, Alert Class, Service, CI, Closed Reason, Closed, CompId, Count, First Occ, ID, Last Occ, Owner, Primary Service, Sup, TicketGroup, TicketID.</b></p> <p>Filters limit display content and drop-down menu selections to only those items that pass through the selected filter's criteria. If no items match the filter, you might have zero search results (an empty table).</p>
<b>Clear</b>	Clears the <b>Field Filter</b> and <b>Search Text</b> entries.
<b>Search Text</b>	Enter the (case-sensitive) string to search for in the selected <b>Field Filter</b> .
<b>CMDB Filter</b>	<p>Shows the selected Owner, Area, Group, Service and Environment filters. By default, all components of the CMDB (*) are included in the search.</p> <p>These <b>CMDB Filter</b> fields are populated when you click Open Alerts Table , which is accessible from the <b>Multi Area Service Views</b> displays, to open the <b>Alerts Table</b> in a new window. The filters selected in the <b>All Management Areas</b> and <b>Multi Area Service Views</b> displays are applied to the <b>Alerts Table</b> (that opens in the new window). NOTE: When you use the navigation tree (in the left panel) to open the <b>Alerts Table</b> display, the <b>Environment</b> filter is applied to the display if it has a value other than * (asterisk).</p>
<b>Clear CMDB Filter</b>	Clears all of the values in the <b>CMDB Filter</b> ( <b>Owner, Area, Group, Service</b> and <b>Environment</b> filters). NOTE: This action is not applied to any other display.
<b>RegEx</b>	Toggles the <b>Search Text</b> field to accept Regular Expressions for filtering.
<b>All</b>	Click to show all alerts in the table: <b>Open</b> and <b>Closed</b> alerts.
<b>Open</b>	Click to only show <b>Open</b> alerts in the table.
<b>Closed</b>	Click to only show <b>Closed</b> alerts in the table.
<b>Owner Filter</b>	Select the alert <b>Owner</b> to show alerts for in the table.
	<p><b>All</b> Shows alerts for all Owners in the table: <b>Not Owned</b> and <b>Owned By Me</b> alerts.</p> <p><b>Not Owned</b> Shows only alerts without Owners in the table.</p> <p><b>Owned By Me</b> Shows only alerts for the current user in the table.</p>
<b>Alert Settings Conn OK</b>	<p>The Alert Server connection state:</p> <p> Disconnected.</p> <p> Connected.</p>
<b>Total</b>	<b>X/Y</b> where <b>X</b> is the total number of alerts in the table with all selected filters applied. <b>Y</b> is the number of alerts in the table with only the <b>CMDB</b> and <b>Cleared</b> filters applied.
<b>Critical</b>	<p>Check to show alerts in the table that are currently in a critical state. NOTE: You must check <b>Critical</b> to see alerts that are in a critical state.</p> <p><b>X/Y</b> where <b>X</b> is the total number of critical alerts in the table with all selected filters applied. <b>Y</b> is the number of alerts in the table with only the <b>CMDB Filter</b> and <b>Cleared</b> filters applied.</p>

<b>Warning</b>	Check to show alerts in the table that are currently in a warning state. NOTE: You must check <b>Warning</b> to see alerts that are in a warning state. <b>X/Y</b> where <b>X</b> is the total number of warning alerts in the table with all selected filters applied. <b>Y</b> is the number of alerts in the table with only the <b>CMDB</b> and <b>Cleared</b> filters applied.
<b>Suppressed</b>	Check to show alerts in the table that are suppressed. The <b>Suppressed</b> count is not impacted by the <b>Critical</b> and <b>Warning</b> filters. It is impacted only by the <b>CMDB Filter</b> and the <b>Owner Filter</b> . NOTE: You must check <b>Suppressed</b> to see Suppressed alerts in the table.
<b>Own</b>	Click to assign an Owner for the alert. This option is only visible when logged in as one of the following roles: event, full, admin, super. This option is disabled when you select a gray row.
<b>Suppress</b>	Click to suppress the alert. This option is only visible when logged in as one of the following roles: event, full, admin, super. This option is disabled when you select a gray row.
<b>UnSuppress</b>	Click to unsuppress the alert. This option is only visible when logged in as one of the following roles: event, full, admin, super. This option is disabled when you select a gray row or when you select a row.
<b>Close</b>	Click to close the alert. This option is only visible to users with Administrator privileges. This option is disabled when you select a gray row or you select a row where the Primary Service is not in the \$rtvManageableCompID list for the logged in user.

### Alerts Table

This table lists all active alerts for the current filters. The table is empty unless you check **Critical**, **Warning**, or both. Filter the list using the search fields and drop-down menus (in the upper portion of the display). To view details about an alert, select an alert and click **Details** (in the bottom right portion of the display) to open the **Alert Detail** dialog. To view details about the CI source of the alert, select an alert and click **Go To CI** (in the bottom right portion of the display) to open its Summary display.

<b>First Occ</b>	The date and time the alert first occurred.
<b>Last Occ</b>	The date and time the alert last occurred.
<b>Count</b>	The number of times the alert was generated.
<b>Sup</b>	When checked, the alert has been suppressed by a user.
<b>Owner</b>	The named owner assigned by the administrator.
<b>Alert Name</b>	The name of the alert.
<b>Primary Service</b>	The name of the Service with which the alert is associated.
<b>CI</b>	The CI alert source.
<b>Alert Text</b>	Description of the alert.
<b>AlertClass</b>	An optional alert field which can be used when integrating with other alerting systems.
<b>CompID</b>	An optional alert field which can be used when integrating with other alerting systems.
<b>TicketID</b>	An optional alert field which can be used when integrating with other alerting systems.
<b>TicketGroup</b>	An optional alert field which can be used when integrating with other alerting systems.

<b>Columns</b>	<p><b>Id</b> When checked, shows the <b>ID</b> column in the table.</p> <p><b>Closed</b> When checked, shows the <b>Closed</b> column in the table.</p> <p><b>Closed Reason</b> When checked, shows the <b>Closed Reason</b> column in the table.</p> <p><b>Alert Index</b> When checked, shows the <b>Alert Index</b> column in the table.</p>
<b>Go To CI</b>	Select an alert from the <b>Alerts Table</b> , then click <b>Go To CI</b> to view details for the selected CI in the Summary display.
<b>Annotate</b>	<p>Select one or more alerts from the <b>Alerts Table</b>, then click <b>Annotate</b> to open the <b>Set Owner and Comments</b> dialog and enter comments or change alert owner. This option is only visible when logged in as one of the following roles: event, full, admin, super. This option is disabled when you select a gray row or when you select a row where the Primary Service is not in the \$rtvManageableCompID list for the logged in user. For details, see <b>Configure User and Role Management</b>.</p> <p><b>ID</b> Lists the alert IDs, separated by semicolons, for the alerts selected from the <b>Alert Table</b>.</p> <p><b>Source</b> Lists the name of the back-end Data Server reporting the alert, separated by semicolons.</p> <p><b>Enter Owner</b> Enter the name of the owner for one or more alerts, click <b>Set Owner of One Alert</b> to assign the Owner, then click <b>Close</b>. By default, this field displays the current user name.</p> <p><b>Enter Comment</b> Enter a comment for one or more alerts, click <b>Add Comment on One Alert</b> to apply the Comment, then click <b>Close</b>. By default, this field displays previously entered comments. The text appears in the <b>Comments</b> field for the alert.</p> <p><b>Set Owner</b> Applies the name of the alert owner in the <b>Enter Owner</b> field for one or more alerts.</p> <p><b>Add Comment</b> Applies the comment in the <b>Enter Comment</b> field for one or more alerts.</p> <p><b>Clear Comments</b> Removes all comments for one or more alerts.</p> <p><b>Close</b> Closes the dialog.</p>
<b>Options</b>	Select a single alert from the <b>Alerts Table</b> , then click <b>Options</b> to open the <b>Alert Options</b> dialog. This dialog is provided for customizing your own alert options. This option is disabled when you select a gray row or more than one row.
<b>Details</b>	Select a single alert from the <b>Alerts Table</b> , then click <b>Details</b> to open the <b>Alert Detail</b> window and view alert details. This option is disabled when you select a gray row or more than one row.

## Alerts Table - HTML

This describes the HTML version of the Alerts Table which can be found under the **Alerts** tab.

Use this display to track and manage all alerts that have occurred in the system, where:



One or more alerts exceeded their ALARM LEVEL threshold in the table row



One or more alerts exceeded their WARNING LEVEL threshold in the table row

You can search, filter, sort and choose columns to include by clicking a column header icon (located to the right of each column label) and selecting **Filter**, **Sort Ascending**, **Sort Descending** or **Columns**. Use the **Ack'd** and **Cleared** drop-downs to filter the table by those columns. Right-click on a table cell to **Export to Excel** or **Copy Cell Value**. Use **Ctrl + click** or **Shift + arrow** to select multiple alerts. To investigate, select one alert and click:

**Details**

to open the **Component Alert Detail** display to get details about that particular alert instance as it specifically applies to the associated CI.

**CI**

to see utilization conditions for the CI associated with the alert during the seconds (minutes, hours or days) leading up to the alert being executed in a summary display.

With one or more alerts selected, you can click **Own** to set the alert(s) owner field, **Ack** to acknowledge the alert(s), **Unack** to clear the acknowledgement on previously acknowledged alert(s) and **Comment** to add a comment to the alert(s).

You must be logged in as `rtvalertmgr` or `rtvadmin` to perform the **Own**, **Ack**, **Unack**, or **Comment** actions. Otherwise, you get an error dialog.

**Alerts Table** 30-Apr-2019 13:47:46 ✓DATA

Ack'd: all Cleared: false Cmdb Filter: \*\*\*\*\* Alert Count: 92

Time	Ack	Clr	Sevl	Alert Name	Alert Text	Owner	ID	Source	Comments	CI
2019-Apr-30 00:04:07			⚠	JvmNotConnected	Server disconnected		1043	RTV-DATA-TIB		win4
2019-Apr-30 01:34:49			⚠	JvmNotConnected	Server disconnected		1009	Z-SIMDATA-1		local
2019-Apr-30 01:34:49			⚠	JvmNotConnected	Server disconnected		1008	Z-SIMDATA-1		local
2019-Apr-30 01:34:49			⚠	JvmNotConnected	Server disconnected		1007	Z-SIMDATA-1		local
2019-Apr-30 01:34:49			⚠	JvmNotConnected	Server disconnected		1006	Z-SIMDATA-1		local
2019-Apr-30 01:34:49			⚠	JvmNotConnected	Server disconnected		1005	Z-SIMDATA-1		local
2019-Apr-30 01:34:49			⚠	JvmNotConnected	Server disconnected		1004	Z-SIMDATA-1		local
2019-Apr-30 01:34:49			⚠	JvmNotConnected	Server disconnected		1003	Z-SIMDATA-1		local
2019-Apr-30 01:34:49			⚠	JvmNotConnected	Server disconnected		1002	Z-SIMDATA-1		local
2019-Apr-30 01:34:49			⚠	JvmNotConnected	Server disconnected		1001	Z-SIMDATA-1		local
2019-Apr-30 01:34:49			⚠	JvmNotConnected	Server disconnected		1000	Z-SIMDATA-1		local
2019-Apr-30 12:01:02			⚠	JvmCpuPercentHigh	High Alert Limit exceed		1064	Z-SIMDATA-1		local
2019-Apr-30 13:44:01			🚨	JvmCpuPercentHigh	High Warning Limit exc		928739	RTV-DATA-KAF		Insta
2019-Apr-30 13:47:04			🚨	JvmCpuPercentHigh	High Warning Limit exc		928747	RTV-DATA-KAF		Insta
2019-Apr-30 01:36:49			🚨	HostCpuPercentHigh	High Warning Limit exc		1010	Z-SIMDATA-1		defa
2019-Apr-30 01:36:49			🚨	HostCpuPercentHigh	High Warning Limit exc		1010	Z-SIMDATA-1		defa
2019-Apr-30 02:05:10			⚠	HostCpuPercentHigh	High Alert Limit exceed		1011	Z-SIMDATA-1		defa

Page 1 of 3 1 - 40 of 92 items

## Alert History Table

Use this display to track the history of any alert that has occurred in your RTView Enterprise system. There is one row in the table for each update to each alert. The table is limited to **20,000** rows. If there are more than **20,000** rows in the selected time range, the newest **20,000** rows are shown.

The color coded navigation tree shows the contents of the CMDB hierarchically ordered. Choose a node to filter alerts shown in the table. The **Alert History Table** only shows alerts associated with the node you select. A green indicator means the node has no associated alerts. A red indicator means the node has one or more associated alerts.

Service name labels are appended with the Environment. For example, the following illustrates that the **TAS-MEMBER** Service currently has no alerts in the **PRODUCTION** Environment.

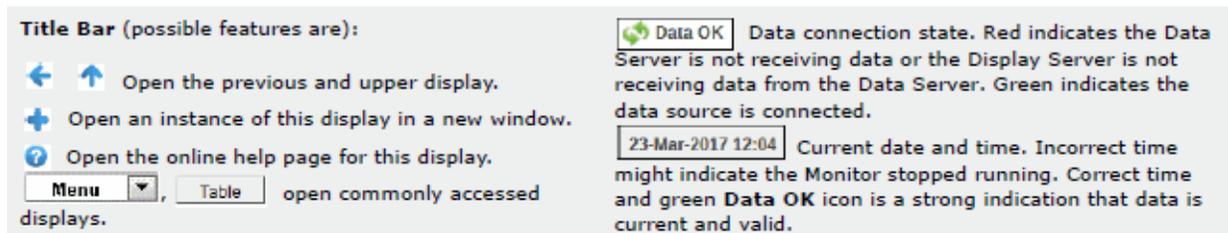


To filter the table, select a table column from the **Field Filter** drop-down menu. In the **Search Text** field, enter the (case-sensitive) string to search for in the selected **Field Filter**, then click **<Enter>**. Click **Clear** to clear the **Field Filter** and **Search Text** fields.

The **Count** label shows two values: the filtered row count / the total row count.

Click a column heading to sort the table by the column data.

The screenshot shows the 'Alert History Table' interface. At the top, there are filters for 'Field Filter' and 'Search Text'. The table has columns: Time, ID, Cif, Sup, Owner, Alert Name, Alert Index, Alert Text, and Cleared Reason. The table contains multiple rows of alert data, with the first few rows highlighted in red. The 'Alert Text' column contains detailed error messages, such as 'High Alert Level exceeded, current value: 104411 8402907147 level: 0000 0'. The interface also shows a 'Count: 3096/3096' label in the top right corner.




---

**Note:** The **Count** field in the title bar of this display shows two values: the filtered row count and the unfiltered row count.

---

The row color indicates the most critical alert state for the row, as follows:

**Row Color Code:**

Tables with colored rows indicate the following:

- Red indicates that one or more alerts exceeded their ALARM LEVEL threshold in the table row.
- Yellow indicates that one or more alerts exceeded their WARNING LEVEL threshold in the table row.
- Green indicates that no alerts exceeded their WARNING or ALARM LEVEL threshold in the table row.

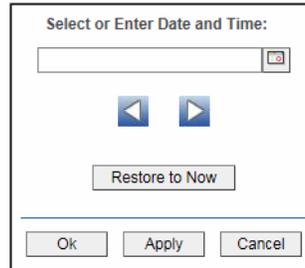
**Fields and Data**

This display includes:

- Field Filter** Select a table column from the drop-down menu to perform a search in: **Alert Name, Alert Text, Cleared Reason, Clr, ID, Owner, Sev, Source, Sup, ID** or **Time**. Filters limit display content and drop-down menu selections to only those items that pass through the selected filter's criteria. If no items match the filter, you might have zero search results (an empty table).
- Clear** Clears entries in the **Alert Name Filter** field and all table data.
- Search Text** Enter the (case-sensitive) string to search for in the selected **Field Filter**.
- RegEx** Toggles the **Search Text** field to accept Regular Expressions for filtering.
- Sort by ID + Time** When checked, table rows are sorted by the **Time** and **ID** columns.

**Time Range**

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu. Click **Restore to Now** to reset the time range end point to the current time.

**Alerts Table**

This table lists all alerts for all Owners and all Areas that have occurred in your RTView Enterprise system. Filter the list by alert names using the **Alert Name Filter** drop-down menu.

<b>Time</b>	The date and time the alert first occurred.
<b>ID</b>	The unique string identifier for the alert.
<b>Clear</b>	When checked, the alert has been cleared by a user.
<b>Sup</b>	When checked, the alert has been suppressed by a user.
<b>Owner</b>	The named owner assigned by the administrator.
<b>Alert Name</b>	The name of the alert.
<b>Alert Index</b>	Lists the Alert Indexes, separated by tildes (~), for the alert.
<b>Alert Text</b>	Descriptive text about the alert.
<b>Cleared Reason</b>	<b>DATA UPDATE:</b> The metric returned to normal thresholds. <b>MANUAL:</b> A user cleared or closed the alert manually.
<b>Sev</b>	The severity level of the alert.
<b>Source</b>	The name of the back-end Data Server reporting the alert.

## Administration

These displays enable you to set alert thresholds, and observe how alerts are managed. Displays in this View are:

- **"Alert Administration"**: Displays active alerts and provides interface to modify and manage alerts.
- **"Tabular Alert Administration"**: Set override alerts.
- **"Alert Administration - HTML"**: The HTML version of the **Alert Administration** display which shows active alerts and provides an interface to modify and manage alerts.
- **"Alert Overrides Administration - HTML"**: The HTML version used for setting override alerts.
- **"Alert Admin Audit"**: Track modifications of alerts throughout your system, such as alert threshold modifications.
- **"Alert Action Audit Trail"**: Track alert management throughout your system, including the name of the user who performed the action, the time the action was performed and what the action was.

## Alert Administration

Set global or override alert thresholds. Alert settings are global by default. Only users logged in with the admin or super roles can save changes to alert thresholds. For details, see **Configure User and Role Management**.

The table describes the global settings for all alerts on the system. To filter the alerts listed in the table, enter a string in the **Alert Filter** field and press **<enter>** or click elsewhere in the display. Filters are case sensitive and no wildcard characters are needed for partial strings. For example, if you enter Server in the **Alert Filter** field, it filters the table to show only alerts with **Server** in the name. Choose **Clear** to clear the filter.

### Global Thresholds

To set a global alert, select an alert from the **Active Alert Table**. The name of the selected alert populates the **Settings for Selected Alert Name** field. Edit the **Settings for Selected Alert** and click **Save Settings** when finished.

The manner in which global alerts are applied depends on the Solution Package. For example, the EMS Monitor Solution Package has queue alerts, topic alerts and server alerts. When a queue alert is applied globally, it is applied to all queues on all servers. Likewise, a server alert applies to all servers, and a topic alert applies to all topics on all servers.

### Override Thresholds

Setting override alerts allows you to set thresholds for a single resource (for example, a single server). Override alerts are useful if the majority of your alerts require the same threshold setting, but there are other alerts that require a different threshold setting. For example, you might not usually be concerned with execution time at a process level, but perhaps certain processes are critical. In this case, you can apply alert thresholds to each process individually.

To apply an individual alert you Index the Monitored Instance or resource (such as a message queue, in the case of the EMS Monitor package). The Index Types available are determined by the Solution Package installed. For example, with the EMS Monitor package you can set an alert for a specific topic on a specific server--the PerServerTopic Index option--rather than for all topics on all servers.

**Note:** To filter the alerts shown in the **Administration - Alert Administration** display by solution package, use the `$rtvAlertPackageMask` substitution. For details, refer to the *RTView Enterprise Configuration Guide*.

The screenshot shows the 'Alert Administration' interface. At the top, there is a title bar with a back arrow, the title 'Alert Administration', the date and time '23-Sep-2015 16:15', and status indicators for 'Data OK' and connection status. Below the title bar is an 'Alert Filter' input field with a 'Clear' button and an 'Alert Settings Conn OK' indicator. The main area contains a table of alerts with columns: Alert, Warning Level, Alarm Level, Duration, Alert Enabled, and Override Count. Below the table is a 'Settings for Selected Alert' section with fields for Name, Description, Warning Level, Alarm Level, Duration (Secs.), and Enabled. There are also 'Save Settings' and 'Override Settings' buttons.

Alert	Warning Level	Alarm Level	Duration	Alert Enabled	Override Count
AcwInstanceCpuHigh	50	75	30	<input checked="" type="checkbox"/>	
AcwInstanceDiskReadBytesHigh	100000	200000	30	<input checked="" type="checkbox"/>	
AcwInstanceDiskReadOpsHigh	100	200	30	<input checked="" type="checkbox"/>	
AcwInstanceDiskWriteBytesHigh	100000	200000	30	<input checked="" type="checkbox"/>	
AcwInstanceDiskWriteOpsHigh	100	200	30	<input checked="" type="checkbox"/>	
AcwInstanceNetworkReadBytesHigh	100000	200000	30	<input checked="" type="checkbox"/>	
AcwInstanceNetworkWriteBytesHigh	100000	200000	30	<input checked="" type="checkbox"/>	
AmxServiceHitRateHigh	200	400	30	<input checked="" type="checkbox"/>	
AmxServiceNodeFaultRateHigh	200	400	30	<input checked="" type="checkbox"/>	
AmxServiceNodeHitRateHigh	200	400	30	<input checked="" type="checkbox"/>	
AmxServiceNodeMovingAvgHitRateHigh	200	400	30	<input checked="" type="checkbox"/>	
AmxServiceNodeMovingAvgResponseTimeHigh	200	400	30	<input checked="" type="checkbox"/>	
AmxServiceNodeResponseTimeHigh	200	400	30	<input checked="" type="checkbox"/>	
AmxServiceResponseTimeHigh	200	400	30	<input checked="" type="checkbox"/>	
Bw6AppNodeCpuUsedHigh	50	80	30	<input type="checkbox"/>	
Bw6AppNodeMemUsedHigh	50	80	30	<input type="checkbox"/>	
Bw6AppProcessCreatedRateHigh	50	80	30	<input checked="" type="checkbox"/>	
Bw6AppProcessElapsedTimeHigh	100	200	30	<input type="checkbox"/>	

**Settings for Selected Alert**

Name:  Warning Level:  Duration (Secs.):

Description:  Alarm Level:  Enabled:

Tabular Alert Options

The Warning Level, Alert Level and Alarm Enabled values on this screen can be overridden for each alert

#### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

#### Fields and Data

This display includes:

- Alert Filter** Enter the (case-sensitive) string to filter the table by the **Alert** table column value. NOTE: Partial strings can be used without wildcard characters. Press **<enter>** or click elsewhere in the display to apply the filter.
- Clear** Clears the **Alert Filter** entry.

**Alert Settings** The Alert Server connection state:

- Disconnected.
- Connected.

### Active Alert Table

This table describes the global settings for all alerts on the system. Select an alert. The name of the selected alert populates the **Settings for Selected Alert Name** field (in the lower panel). Edit **Settings for Selected Alert** fields and click **Save Settings** when finished.

<b>Alert</b>	The name of the alert.
<b>Warning Level</b>	The global warning threshold for the selected alert. When the specified value is exceeded a warning is executed.
<b>Alarm Level</b>	The global alarm threshold for the selected alert. When the specified value is exceeded an alarm is executed.
<b>Duration (Secs)</b>	The amount of time (in seconds) that the value must be above the specified Warning Level or Alarm Level threshold before an alert is executed. <b>0</b> is for immediate execution.
<b>Alert Enabled</b>	When checked, the alert is enabled globally.
<b>Override Count</b>	The number of times thresholds for this alert have been defined individually in the <b>Tabular Alert Administration</b> display. A value of: <b>-0</b> indicates that no overrides are applied to the alert. <b>-1</b> indicates that the alert does not support overrides.

### Settings for Selected Alert

To view or edit Global settings, select an alert from the **Active Alert Table**. Edit the **Settings for Selected Alert** fields and click **Save Settings** when finished.

To set override alerts, click on **Override Settings** to open the **Tabular Alert Administration** display.

<b>Name</b>	The name of the alert selected in the <b>Active Alert Table</b> .
<b>Description</b>	Description of the selected alert. Click Calendar <input type="text"/> for more detail.
<b>Warning Level</b>	Set the Global warning threshold for the selected alert. When the specified value is exceeded a warning is executed. To set the warning to occur sooner, reduce the Warning Level value. To set the warning to occur later, increase the Warning Level value. NOTE: For low value-based alerts (such as <b>EmsQueuesConsumerCountLow</b> ), to set the warning to occur sooner, increase the Warning Level value. To set the warning to occur later, reduce the Warning Level value.
<b>Alarm Level</b>	Set the Global alarm threshold for the selected alert. When the specified value is exceeded an alarm is executed. To set the alarm to occur sooner, reduce the Alarm Level value. To set the warning to occur later, increase the Alarm Level value. NOTE: For low value-based alerts (such as <b>EmsQueuesConsumerCountLow</b> ), to set the alarm to occur sooner, increase the Alarm Level value. To set the alarm to occur later, reduce the Alarm Level value.
<b>Duration</b>	Set the amount of time (in seconds) that the value must be above the specified Warning Level or Alarm Level threshold before an alert is executed. <b>0</b> is for immediate execution. This setting is global.
<b>Enabled</b>	Check to enable alert globally.

**Save Settings** Click to apply alert settings.

**Override Settings** Click to open the **Tabular Alert Administration** display to set override alerts on the selected alert.

## Tabular Alert Administration

Set override alerts (override global alert settings). This display opens when you select an alert in the **Alert Administration** display and then select **Override Settings**.

For step-by-step instructions setting thresholds for individual alerts, see **Setting Override Alerts..**

### Fields and Data

This display includes:

- Alert Settings Conn OK** The connection state.
- No servers are found.
  - One or more servers are delivering data.

### Override Settings For Alert:(name)

This table lists and describes alerts that have override settings for the selected alert. Select a row to edit alert thresholds. The selected item appears in the Index field. Edit settings in the Alert Settings fields, then click Save Settings.

<b>Index Type</b>	Select the type of alert index to show in the Values table. Options in this drop-down menu are populated by the type of alert selected, which are determined by the Package installed. For example, with the EMS Monitor package the following Index Types are available: <ul style="list-style-type: none"> <li>• PerServer: Alert settings are applied to a specific server.</li> <li>• PerQueue: Alert settings are applied to the queue on each server that has the queue defined.</li> <li>• PerServerQueue: Alert settings are applied to a single queue on a specific server.</li> <li>• PerTopic: Alert settings are applied to the topic on each server that has the topic defined.</li> <li>• PerServerTopic: Alert settings are applied to a single topic on a specific server.</li> </ul>
<b>Index</b>	The value of the index column.
<b>Override Settings</b>	When checked, the override settings are applied.
<b>Alert Enabled</b>	When checked, the alert is enabled.
<b>Index Type</b>	Select the index type. The index type specifies how to apply alert settings. For example, to a queue (topic or JVM, and so forth) across all servers, or to a queue on a single server. NOTE: Options in this drop-down menu are populated by the type of alert selected from the Alert Administration display. Index Types available depend on the Package installed.
<b>Index</b>	The selected index column to be edited. This field is populated by the selection made in the <b>Unassigned Indexes</b> table.
<b>Unassigned Indexes</b>	This table lists all possible indexes corresponding to the Index Type chosen in the drop-down list. Select a row to apply individual alert thresholds. The selected item appears in the <b>Index</b> field. Edit settings in the <b>Alert Settings</b> fields, then click <b>Add</b> .
<b>Add</b>	Click to add changes made in <b>Alert Settings</b> , then click <b>OK</b> to confirm.
<b>Remove</b>	Click to remove an alert selected in the <b>Index Alert Settings</b> table, then click <b>OK</b> to confirm.
<b>Save Settings</b>	Click to save changes made to alert settings.
<b>Alert Settings</b>	Select a topic, server or queue from the <b>Unassigned Indexes</b> table and edit the following settings.
<b>Warning Level</b>	Set the warning threshold for the selected alert. When the specified value is exceeded a warning is executed. To set the warning to occur sooner, reduce the Warning Level value. To set the warning to occur later, increase the Warning Level value. NOTE: For low value-based alerts (such as <b>EmsQueuesConsumerCountLow</b> ), to set the warning to occur sooner, increase the Warning Level value. To set the warning to occur later, reduce the Warning Level value. <b>Click Save Settings to save settings.</b>

<b>Alarm Level</b>	Set the alarm threshold for the selected alert. When the specified value is exceeded an alarm is executed. To set the alarm to occur sooner, reduce the Alarm Level value. To set the warning to occur later, increase the Alarm Level value. NOTE: For low value-based alerts (such as <b>EmsQueuesConsumerCountLow</b> ), to set the alarm to occur sooner, increase the Alarm Level value. To set the alarm to occur later, reduce the Alarm Level value. Click <b>Save Settings</b> to save settings.
<b>Alert Enabled</b>	Check to enable the alert, then click <b>Save Settings</b> .
<b>Override Settings</b>	Check to enable override global setting, then click <b>Save Settings</b> .

**Back to Alerts** Returns to the **Administration - Alert Administration** display.

### Setting Override Alerts

Perform the following steps to set an override alert. Index Types available depend on the Solution Package installed. In this example, we use the EMS Monitor Package to illustrate.

---

**Note:** To turn on an alert, both **Alert Enabled** and **Levels Enabled** must be selected.

---

To turn on/off, change threshold settings, enable/disable or remove an alert on a single resource:

1. In the **Alert Administration** display, select a tabular alert in the **Active Alert Table** and click **Override Settings**. The **Tabular Alert Administration** display opens.

---

**Note:** Alerts that do not support overrides have a value of **-1** for the **Override Count** column and the **Override Settings** option is not present when you select such an alert.

---

2. In the **Tabular Alert Administration** display, select the Index type from the **Index Type** drop-down menu (options are populated by the type of alert you previously selected). For example, with the EMS Monitor package, select PerServerQueue, PerServerTopic or PerServer. NOTE: If you select PerServerQueue or PerServerTopic, the alert settings are applied to the queue or topic on a single server.
3. In the **Unassigned Indexes** table, select the item you want to apply an override alert setting to, click **Add** and **OK** in the confirmation dialog. After a few moments the override setting appears in the **AlertLevels** table.
4. Select the item in the **AlertLevels** table.
5. In the Alert Settings panel (lower right), if needed, modify the Warning Level and Alarm Level settings.
6. In the **Alert Settings** panel, set the following as appropriate.
  - To turn on the alert for this index with the given thresholds:
    - Alert Enabled** Select this option.

**Override Settings** Select this option.

**NOTE:** To turn on an alert, both **Alert Enabled** and **Override Settings** must be selected.

- To turn off the alert for only this index (global alert thresholds will no longer apply to this index):

**Alert Enabled** Deselect this option.

**Override Settings** Select this option.

- To no longer evaluate this indexed alert and revert to global settings (or, optionally, Remove it if it is never to be used again):

**Alert Enabled** Not used.

**Override Settings** Deselect this option.

7. Click **Save Settings**. In a few moments the modifications are updated and a new record appears in the **AlertLevels** table. For example, in the following figure, the EmsServerConnectionCountHigh alert has a new override applied. New overrides increment the alert **Override Count** in the **ALERTLEVELS** table.

Alert	Warning Level	Alarm Level	Duration	Alert Enabled	Override Count
EmsQueuesProducerCountHigh	60	80	30	<input type="checkbox"/>	0
EmsQueuesProducerCountLow	15	5	30	<input type="checkbox"/>	0
EmsServerAsyncDBSizeHigh	50	100	30	<input type="checkbox"/>	0
EmsServerConnectionCountHigh	60	80	30	<input checked="" type="checkbox"/>	1
EmsServerInMsgRateHigh	60	80	30	<input type="checkbox"/>	0
EmsServerMemUsedHigh	60	80	30	<input type="checkbox"/>	0

## Alert Administration - HTML

This describes the HTML version of the Alert Administration display which can be found under the **Admin** tab.

The **Alert Administration** display allows administrators to enable/disable alerts and manage alert thresholds. The table describes the global settings for all alerts on the system.

You can set the **Delay** time (the number of seconds that must pass before an alert is triggered, where **0** sets it to immediately execute).

You can set the **Warning Level** which executes a single warning alert when the number of seconds specified here is exceeded. To set the warning to occur sooner, reduce the **Warning Level** value. To set the warning to occur later, increase the **Warning Level** value.

You can set the **Alarm Level** which executes a single alarm alert when the number of seconds specified here is exceeded. To set the alarm to occur sooner, reduce the **Alarm Level** value. To set the alarm to occur later, increase the **Alarm Level** value.

**Note:** For low value-based alerts (an alert that executes based on a value going below a certain threshold), to set the alarm to occur sooner you increase the **Alarm Level** value. To set the alarm to occur later, reduce the **Alarm Level** value.

You can apply alert thresholds globally or as an *override*. Setting override alerts allows you to set thresholds for a subset of your resources, or for a single resource (for example, a single server). Override alerts are useful if the majority of your resources require the same threshold setting, but there are a few resources that require a different threshold setting. For example, you might not usually be concerned with execution time at a process level, but perhaps certain processes are critical. In this case, you can apply alert thresholds to each process individually. See below for instructions.

You can filter, sort and choose columns to include by clicking a column header icon (located to the right of each column label) and selecting **Filter**, **Sort Ascending**, **Sort Descending** or **Columns**. Use the **Ack'd** and **Cleared** drop-downs to filter the table by those columns. Right-click on a table cell to **Export to Excel**.

### To set thresholds and enable a global alert:

Select an alert and, under **Settings for alert** (in the lower portion of the screen), modify settings for the alert **Delay**, **Warning Level** and/or **Alarm Level** and **Save Settings**. With that alert selected, check the **Alert Enabled** box under **Settings for alert** (in the lower portion of the screen) and **Save Settings**. The **Alert Enabled** box (next to the selected alert) is now checked.

### To set thresholds and enable an override alert:

To set an override alert, select an alert and click **Override Settings** to open the **Alert Overrides Admin** display.

**Alerts Administration**
30-Apr-2019 10:34:01 ✓ DATA OK

Package: All [http://rtvdemos.sl.com/emdemo\\_central\\_rtvquery](http://rtvdemos.sl.com/emdemo_central_rtvquery)

Alert Name	Alert Enabled	Alert Delay	Warning Level	Alert Level	Override Count
HostNetworkTxRateHigh	<input type="checkbox"/>	30	50	75	0
HostProcessCountLow	<input type="checkbox"/>	30	15	5	0
HostStateData	<input type="checkbox"/>	30			0
HostStorageUsedHigh	<input type="checkbox"/>	30	80	90	0
HostSwapUsedHigh	<input type="checkbox"/>	30	75	90	0
HostVirtualMemoryUsedHigh	<input type="checkbox"/>	30	75	90	0
JvmCpuPercentHigh	<input checked="" type="checkbox"/>	60	50	70	0
JvmGcDutyCycleHigh	<input type="checkbox"/>	30	50	75	0
JvmMemoryUsedAfterGCHigh	<input type="checkbox"/>	0	1	80	0
JvmMemoryUsedHigh	<input checked="" type="checkbox"/>	60	75	86	0
JvmNotConnected	<input checked="" type="checkbox"/>	60			0
JvmStateData	<input type="checkbox"/>	30			0
JvmThreadCountHigh	<input checked="" type="checkbox"/>	60	8000	12000	0

Page 2 of 5 101 - 200 of 432 items

**Settings for alert**

Alert Enabled:       Delay:       Warning Level:       Alert Level:

Alert Selected: **HostSwapUsedHigh**    Description: **The percentage of swap space used is above the limits defined for that Host**

See [“Alert Overrides Administration - HTML”](#) for additional details.

## Alert Overrides Administration - HTML

This describes the HTML version of the Alert Overrides Administration display.

Administrators use this display to create override alerts. To access this display, select an alert in the **Alert Administration** display and choose **Override Settings**.

The table lists all the resources to which you can apply the alert you selected from the **Alert Administration** display. Each row in the table is a different resource, columns describe whether that alert is enabled (globally and as an override) and if so, the current alert thresholds for each.

You can filter, sort and choose columns to include by clicking a column header icon (located to the right of each column label) and selecting **Filter**, **Sort Ascending**, **Sort Descending** or **Columns**. Use the **Display** drop-down to filter the table to show **All** resources, only resources with the **Overridden** alert applied or **Free** resources (to show only resources without the alert override applied). Right-click on a table cell to **Export to Excel** or **Copy Cell Value**.

### To set an override alert:

Select a resource and **Override Type** from the drop-down menu (depending on the alert, there might be only one type). Under **Settings for selected index** (in the lower portion of the screen), modify settings for the alert **Delay**, **Warning Level** and/or **Alarm Level** and **Add Override**. The table updates with your new settings.

With that resource selected, toggle on  **Override Enabled** and **Alert Enabled** under **Settings for alert** (in the lower portion of the screen) and **Save Settings**. The table updates with your new settings. The **Alert Administration** display **Override Count** also updates for the alert.

**Alert Overrides Administration** Data Server: RTV-DATA-INFRA 20-Apr-2019 15:48:05 DATA OK

Alert: HostCpuLoadAvg1High    Override Type: PerHost    Display: All

domain	hostname	Override Enabled	Alert Enabled	Warning Level	Alert Level
SL-DEMO-LX	192.168.200.201				
SL-DEMO	SLHOST13				
SL-DEMO	SLHOST14				
SL-DEMO	SLHOST3				
SL-DEMO-LX	192.168.200.42				
SL-DEMO	SLHOST20				
SL-DEMO-LX	192.168.200.92				
SL-DEMO-LX	192.168.200.91				
SL-DEMO	SLHOST93				
SL-DEMO	SLHOST1				
SL-DEMO	SLHOST10				
SL-DEMO	SLRTVMGR				
SL-DEMO	SLHOST2				
SL-DEMO-LX	192.168.200.89	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	60	85
SL-DEMO	vmxp-16				

**Settings for selected index**

Override Enabled:     Alert Enabled:     Warning Level:     Alert Level:

## Alert Admin Audit

View alert management details such as alert threshold modifications.

Each table row is a single modification made to an alert. To view modifications for a single alert in a group, click Sort  to order the **ALERTNAME** column.

Alert Administration Audit Trail							23-Sep-2015 16:08	Data OK	+	?
Audit Conn OK										
TIME_STAMP	USER	ACTION	ALERTNAME	INDEXTYPE	ALERTINDEX	WARNINGLEVE				
09/20/15 15:27:45	admin	UPDATED	BwActivityErrorRateHigh	Default	Default	0.0				
09/20/15 15:16:15	admin	UPDATED	BwActivityExecutionTimeHigh	Default	Default	0.0				
09/20/15 15:16:00	admin	UPDATED	BwActivityErrorRateHigh	Default	Default	0.0				
09/19/15 10:35:32	admin	UPDATED	BwProcessElapsedTimeHigh	Default	Default	0.0				
09/19/15 10:35:20	admin	UPDATED	BwProcessElapsedTimeHigh	Default	Default	0.0				
09/19/15 10:35:07	admin	UPDATED	BwProcessAbortRateHigh	Default	Default	0.0				
09/19/15 10:34:56	admin	UPDATED	BwProcessAbortRateHigh	Default	Default	0.0				
09/19/15 10:34:43	admin	UPDATED	BwEngineCpuUsedHigh	Default	Default	0.0				
09/19/15 10:34:32	admin	UPDATED	BwEngineCpuUsedHigh	Default	Default	0.0				
09/19/15 10:34:12	admin	UPDATED	BwEngineMemUsedHigh	Default	Default	0.0				
09/19/15 10:34:00	admin	UPDATED	BwEngineMemUsedHigh	Default	Default	0.0				
09/19/15 10:33:47	admin	UPDATED	BwEngineCpuUsedHigh	Default	Default	0.0				
09/19/15 10:33:36	admin	UPDATED	BwEngineCpuUsedHigh	Default	Default	0.0				
09/19/15 10:33:21	admin	UPDATED	BwActivityExecutionTimeHigh	Default	Default	0.0				
09/19/15 10:33:06	admin	UPDATED	BwActivityExecutionTimeHigh	Default	Default	0.0				
09/19/15 10:32:50	admin	UPDATED	BwActivityErrorRateHigh	Default	Default	0.0				
09/19/15 10:32:19	admin	UPDATED	BwActivityErrorRateHigh	Default	Default	0.0				
09/19/15 09:42:07	admin	UPDATED	BwEngineCpuUsedHigh	Default	Default	0.0				
09/19/15 09:41:42	admin	UPDATED	BwActivityExecutionTimeHigh	Default	Default	0.0				
09/19/15 09:41:30	admin	UPDATED	BwActivityExecutionTimeHigh	Default	Default	0.0				
09/19/15 09:40:59	admin	UPDATED	BwActivityErrorRateHigh	Default	Default	0.0				
09/19/15 09:40:30	admin	UPDATED	BwActivityErrorRateHigh	Default	Default	0.0				
09/19/15 09:39:30	admin	UPDATED	BwActivityExecutionTimeHigh	Default	Default	0.0				
09/19/15 09:39:09	admin	UPDATED	BwActivityExecutionTimeHigh	Default	Default	0.0				
09/19/15 09:34:23	admin	UPDATED	BwActivityExecutionTimeHigh	Default	Default	0.0				
09/19/15 09:34:07	admin	UPDATED	BwActivityErrorRateHigh	Default	Default	0.0				

#### Title Bar (possible features are):

-   Open the previous and upper display.
-  Open an instance of this display in a new window.
-  Open the online help page for this display.
- open commonly accessed displays.

 Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

#### Fields and Data

This display includes:

- Audit Conn OK** The Alert Server connection state.
  -  Disconnected.
  -  Connected.
- TIME\_STAMP** The date and time of the modification.
- USER** The user name of the administrator who made the modification.
- ACTION** The type of modification made to the alert, such as **UPDATED**.
- ALERTNAME** The name of the alert modified.
- INDEXTYPE** The type of alert Index.
 

Index Type refers to the manner in which alert settings are applied and vary among CI Types. For example, the JVM CI Type has a PerJvm Index Type, the EMS CI Type has PerServer, PerTopic and PerQueue Index Types which apply alerts to servers, topics and queues, respectively.

- ALERTINDEX** The index of the alert which identifies its source.
- WARNINGLEVEL** The warning threshold value for the alert at the time this modification was made, as indicated in the **TIME\_STAMP** column.  
The warning level is a threshold that, when exceeded, a warning is executed.
- ALARMLEVEL** The alarm threshold value for the alert at the time this modification was made, as indicated in the **TIME\_STAMP** column.  
The alarm level is a threshold that, when exceeded, an alarm is executed.
- DURATION** The duration value for the alert at the time this modification was made, as indicated in the **TIME\_STAMP** column.  
The alert duration is the amount of time (in seconds) that a value must be above the specified Warning Level or Alarm Level threshold before an alert is executed. **0** is for immediate execution.
- ENABLED** When checked, indicates the alert was enabled at the time this modification was made, as indicated in the **TIME\_STAMP** column.
- USEINDEX** When checked, indicates the alert override was enabled at the time this modification was made, as indicated in the **TIME\_STAMP** column. For details about alert overrides, see **Alert Administration**.

## Alert Action Audit Trail

The **Alert Action Audit Trail** display shows all user actions concerning alert management, including the name of the user who performed the action, the time the action was performed and what the action was. This display can help managers of the RTView Enterprise solution determine how and when user interactions have impacted the alert system and help manage users so that best practices for alert handling are maintained.

The screenshot shows the 'Alert Action Audit Trail' window. At the top right, it displays the date and time '13-Oct-2015 11:03' and a 'Data OK' status with a green icon. Below the title bar is a table with the following columns: TIME\_STAMP, USER, ACTION\_TYPE, ACTION, TARGET, VALUE, and ALERT\_NAME. The table contains 28 rows of data, all performed by 'admin' on '10/01/15 16:56:29'. The actions are categorized into 'Event Management' (Set Owner) and 'Clear Alert'. The targets include various system components like 'EmsServerRouteState', 'EmsQueueProviderIdleT', 'EmsTopicsProducerCou', and 'EmsTopicsConsumerCo'.

TIME_STAMP	USER	ACTION_TYPE	ACTION	TARGET	VALUE	ALERT_NAME
10/01/15 16:56:29	admin	Event Management	Set Owner	2764	admin	EmsServerRouteState
10/01/15 16:56:29	admin	Event Management	Set Owner	2562	admin	EmsQueueProviderIdleT
10/01/15 16:56:29	admin	Event Management	Set Owner	2385	admin	EmsQueueProviderIdleT
10/01/15 16:56:29	admin	Event Management	Set Owner	2339	admin	EmsTopicsProducerCou
10/01/15 16:56:29	admin	Event Management	Set Owner	2304	admin	EmsTopicsProducerCou
10/01/15 16:56:29	admin	Event Management	Set Owner	2256	admin	EmsTopicsProducerCou
10/01/15 16:56:29	admin	Event Management	Set Owner	2096	admin	EmsTopicsProducerCou
10/01/15 16:56:29	admin	Event Management	Set Owner	2039	admin	EmsTopicsConsumerCo
10/01/15 16:56:29	admin	Event Management	Set Owner	2004	admin	EmsTopicsConsumerCo
10/01/15 16:56:29	admin	Event Management	Set Owner	1956	admin	EmsTopicsConsumerCo
10/01/15 16:56:29	admin	Event Management	Set Owner	1796	admin	EmsTopicsConsumerCo
10/01/15 16:56:29	admin	Event Management	Set Owner	1761	admin	EmsServerAsyncDBSize
10/01/15 16:56:29	admin	Event Management	Set Owner	1732	admin	EmsQueuesProducerCo
10/01/15 16:56:29	admin	Event Management	Set Owner	1375	admin	EmsQueuesProducerCo
10/01/15 16:56:29	admin	Event Management	Set Owner	1358	admin	EmsQueuesConsumerC
10/01/15 16:56:29	admin	Event Management	Set Owner	1001	admin	EmsQueuesConsumerC
10/01/15 16:56:29	admin	Event Management	Clear Alert	2764		EmsServerRouteState
10/01/15 16:56:29	admin	Event Management	Clear Alert	2562		EmsQueueProviderIdleT
10/01/15 16:56:29	admin	Event Management	Clear Alert	2385		EmsQueueProviderIdleT
10/01/15 16:56:29	admin	Event Management	Clear Alert	2339		EmsTopicsProducerCou
10/01/15 16:56:29	admin	Event Management	Clear Alert	2304		EmsTopicsProducerCou
10/01/15 16:56:29	admin	Event Management	Clear Alert	2256		EmsTopicsProducerCou
10/01/15 16:56:29	admin	Event Management	Clear Alert	2096		EmsTopicsProducerCou
10/01/15 16:56:29	admin	Event Management	Clear Alert	2039		EmsTopicsConsumerCo
10/01/15 16:56:29	admin	Event Management	Clear Alert	2004		EmsTopicsConsumerCo
10/01/15 16:56:29	admin	Event Management	Clear Alert	1956		EmsTopicsConsumerCo

### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu , Table open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Fields and Data**

This display includes:

<b>Action Audit Conn OK</b>	The Alert Action database connection availability.  The connection to the Alert Action database is not available.  The connection to the Alert Action database is available.
<b>Time_Stamp</b>	The time the action was performed.
<b>User</b>	The log in name of the user who performed the action.
<b>Action_Type</b>	The type of action which was performed.
<b>Action</b>	The action which was performed.
<b>Target</b>	The alert ID on which the action was performed.
<b>Value</b>	Any value associated with the alert action.
<b>Alert_Name</b>	The name of the alert on which the action was performed.
<b>Alert_Index</b>	The index of the alert which identifies its source.

## CMDB Administration

This display allows you to modify your Service Data Model.

- ["CMDB Admin"](#): View or modify your Service Data Model.
- ["CMDB Administration - HTML"](#): Describes the HTML version of this display which allows you to modify your Service Data Model.

### CMDB Admin

Use this display to setup, view or modify your Service Data Model (CMDB), including: adding, renaming, deleting or merging your CMDB hierarchical elements (Owners, Areas, Groups or Services), associating CIs with Services and assigning or modifying CI attributes (such as Criticality). Only users logged in with the admin or super roles can apply changes in this display.

This display requires administrator permissions.

The **CI List for Selected Service** (upper) table lists the CIs that are associated with the Service selected (from the **Service** drop-down menu).

The **Available Components** (lower) table is not part of the CMDB. The **Available Components** table lists all available CIs for the CI Type (selected from the **Selected CI Type** drop-down menu) that are in your RTView Enterprise system--whether or not they are in the CMDB. Filter this list using the **CIName Filter** field.

You add CIs to the CMDB by associating them with an Owner, Area, Group, and Service. To do so, select the CI Type from the **Selected CI Type** drop-down menu, choose one or more CIs from the **Available Components** table, then click **Add CI**.

It is not necessary to restart the Configuration Server after making changes to the Service Data Model using the **CMDB Admin** display.

**Creating a new Service, Group, Area or Owner:**

Select the CI Type from the **Selected CI Type** drop-down menu, choose one or more CIs from the **Available Components** table, then click **Add CI To...** Assign a new or existing Owner, Area, Group or Service, review your entries and click **OK**. Your changes are visible in drop-down menus and displays.

**Associating CIs with a Service:**

This option is useful when you want to define which CIs are to be monitored for Services. CIs can be associated with more than one Service, Group, Area or Owner. Select the Owner, Area, Group and Service to which you want to associate one or more CIs using the drop-down menus. The **CI List Table** (the upper table) populates with all CIs already associated with the Owner, Area, Group and Service you select. Select the CI Type of the CI(s) you want to associate. The **Available Components** table (the lower table) populates with all CIs that are categorized as that CI Type. Select one or more CIs in the **Available Components** table, set the Criticality and other optional assignments using the drop-down menus (on the right). Click **Add CI** to associate the CI(s) with the Service. A row is added for each associated CI to the **CI List Table**. Your changes are visible in the drop-down menus and displays.

**Renaming a Service, Group, Area or Owner:**

This option is useful when, for example, a *new* Owner is replacing a retiring Owner, a name is misspelled or a more relevant name is required. Select the relevant Owner, Area, Group or Service using the drop-down menus, then click the corresponding **Manage (Owner, Area, Group or Service)** option for what you are renaming. The **Manage (Owner, Area, Group or Service)** dialog opens. In the **Manage (Owner, Area, Group or Service)** dialog, type the new name in the **New Name** field, click **Rename** and **OK**. Your changes are visible in the drop-down menus and displays.

**Deleting a Service, Group, Area or Owner:**

This option is useful when, for example, an Owner, Area, Group or Service and all the CIs associated with it are not relevant in your RTView Enterprise system. When you delete a Service, Group, Area or Owner everything underneath it (lower CMDB levels and associated CIs) is also removed from the CMDB database and displays. Select the relevant Owner, Area, Group or Service using the drop-down menus, then click the corresponding **Manage (Owner, Area, Group or Service)** option for what you are deleting. The **Manage (Owner, Area, Group or Service)** dialog opens. In the **Manage (Owner, Area, Group or Service)** dialog click **Delete** and **OK**. Your changes are visible in the drop-down menus and displays.

---

**Important:** There is no option to undo a deletion from the CMDB. To restore a deletion you must recreate the Owner, Area, Group or Service and the CIs must be re-associated.

---

**Moving a Service, Group or Area:**

This option is useful when, for example, an Area belongs under a different Owner, a Group belongs under a different Area or a Service belongs under a different Group. When you move a Service, Group or Area (Owners cannot be moved) everything underneath it (lower CMDB levels and associated CIs) moves with it. Select the Area, Group or Service you want to move using the drop-down menus, then click the relevant **Manage (Area, Group or Service)** option for what you are moving. The **Manage (Area, Group or Service)** dialog opens. In the **Manage (Area, Group or Service)** dialog, select the new Owner, Area, Group or Service to move to from the **New (Area, Group or Service)** drop-down menus, click **Move** and **OK**. Your changes are visible in the drop-down menus and displays.

### Merging Services, Groups, Areas or Owners:

This option is useful when, for example, an *existing* Owner is taking over for a retiring Owner. When you merge a Service, Group, Area or Owner its name changes to that of the target Service, Group, Area or Owner, and everything underneath it (lower CMDB levels and associated CIs) goes with it. Select the Area, Group or Service you want to merge using the drop-down menus, then click the relevant **Manage (Area, Group or Service)** option for what you are merging. The **Manage (Area, Group or Service)** dialog opens. In the **Manage (Area, Group or Service)** dialog, select an existing Owner, Area, Group or Service to merge to in the **New Name** field, click **Merge** and **OK**. Your changes are visible in the drop-down menus and displays.

### Deleting a CI:

Select a CI from the **CI List Table**, click **Delete** and **OK**. The CI is removed from the CMDB database and displays. Your changes are visible in the drop-down menus and displays.

### Applying Criticality value to multiple CIs:

In the **CI List Table** select a CI that has the Criticality value you want to apply to all CIs in the **CI List Table**, click **Update Criticality like selected CI** and **OK**. The **Criticality** column for all CIs is updated. Your changes are visible in the drop-down menus and displays.

### Changing CI attributes

In the **CI List Table** select the CI you want to modify attributes for, use the **Environment**, **Region**, **SiteName**, **Criticality**, **City**, **Country** and **OSType** drop-down menus to apply attributes, then click **Update** and **OK**. The **CI List Table** is updated. Your changes are visible in the drop-down menus and displays.

By default, the Owner named **Infrastructure** is created. **Infrastructure** organizes all available CIs collected through all Data Servers configured under RTView EM by technology. This default organization can be disabled if needed.

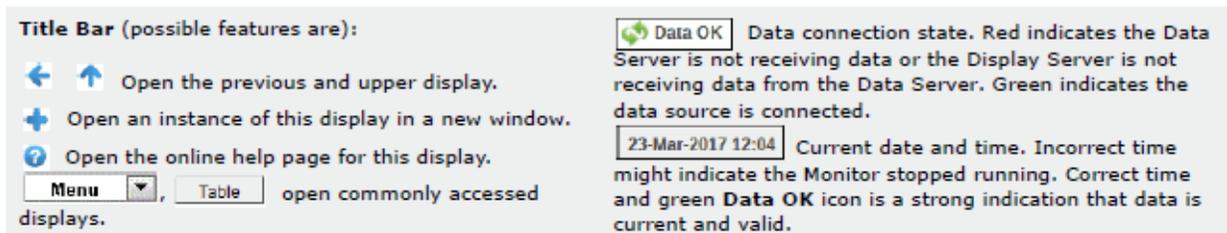
The screenshot shows the 'CMDB - Administration' interface. At the top, it displays the date '24-Sep-2015 11:45' and 'Data OK'. Below this, there are dropdown menus for 'Owner: Jerelyn Parker', 'Area: Systems', 'Group: Databases', and 'Service: IBM DB2'. To the right of these are buttons for 'Manage Owner', 'Manage Area', 'Manage Group', and 'Manage Service'. Further right, there is a 'Source: RTV\_CMDB' label and a button for 'Update Criticality like selected CI'. Below these are several more dropdown menus: 'Environ: DEMOSITE', 'Region:', 'Criticality: A', 'SiteName: Headquarters', 'City:', 'Country:', and 'OSType:'. At the bottom of this section are 'Update' and 'Delete' buttons.

In the center, there is a table titled 'CI List for Selected Service - select a CI to see detail and to edit:'. The table has columns for 'CIType', 'CIName', 'Criticality', 'Region', and 'Env'. It contains two rows of data:

CIType	CIName	Criticality	Region	Env
VMWARE-HOST	vSphere2;slesxi-1_sldemos-hq.local	B		QA
VMWARE-VM	vSphere2;2008S-WIN14	B		QA

Below the table, there is a 'Selected CI Type: MQ-QUEUE' dropdown and a 'CIName Filter:' text input. To the right of the filter is a 'Regex' checkbox and buttons for 'Add CI' and 'Add CI To...'. At the bottom, there is another table titled 'Available Components (CIs):' with columns for 'conn', 'Name', 'CIName', and 'Data Servi'. It lists five rows of test queue components:

conn	Name	CIName	Data Servi
vmrh5-1	TEST_Q_01	vmrh5-1;TEST_Q_01	MQMON-84-OL7-5
vmrh5-1	TEST_Q_02	vmrh5-1;TEST_Q_02	MQMON-84-OL7-5
vmrh5-1	TEST_Q_03	vmrh5-1;TEST_Q_03	MQMON-84-OL7-5
vmrh5-1	TEST_Q_04	vmrh5-1;TEST_Q_04	MQMON-84-OL7-5
vmrh5-1	TEST_Q_05	vmrh5-1;TEST_Q_05	MQMON-84-OL7-5



## Fields and Data

This display includes:

<b>Owner</b>	Select an Owner to filter by. The Owner selected populates the <b>Area</b> , <b>Group</b> and <b>Service</b> drop-down menus.
<b>Manage Owner</b>	<p>Opens a dialog that enables you to <b>Delete</b>, <b>Rename</b> or <b>Merge</b> the Owner.</p> <p><b>Delete</b> removes the Owner from the CMDB database as well as all CMDB data and CIs associated with the Owner.</p> <p><b>Rename</b> Changes all records for the Owner to a new name. <b>Rename</b> is disabled when the name you are typing in the text box already exists in the CMDB.</p> <p><b>Merge</b> Changes all records for the Owner to a different, already existing name in the CMDB. <b>Merge</b> is enabled when the name you are typing in the text box already exists in the CMDB.</p> <p><b>Note:</b> You cannot move Owners.</p>
<b>Area</b>	Select an Area to filter by. The Area selected populates the <b>Group</b> and <b>Service</b> drop-down menus.
<b>Manage Area</b>	<p>Opens a dialog that enables you to <b>Delete</b>, <b>Rename</b> or <b>Merge</b> the Area.</p> <p><b>Delete</b> removes the Area from the CMDB database as well as all CMDB data and CIs associated with the Area.</p> <p><b>Rename</b> Changes all records for the Area to a new name. <b>Rename</b> is disabled when the name you are typing in the text box already exists in the CMDB.</p> <p><b>Merge</b> Changes all records for the Area to a different, already existing name in the CMDB. <b>Merge</b> is enabled when the name you are typing in the text box already exists in the CMDB.</p> <p><b>Move</b> Changes all records for the Area to a different, already existing name in the CMDB that you choose from the <b>New Area</b> drop-down menu.</p>
<b>Group</b>	Select a Group to filter by. The Group selected populates the <b>Service</b> drop-down menu.
<b>Manage Group</b>	<p>Opens a dialog that enables you to <b>Delete</b>, <b>Rename</b> or <b>Merge</b> the Group.</p> <p><b>Delete</b> removes the Group from the CMDB database as well as all CMDB data and CIs associated with the Group.</p> <p><b>Rename</b> Changes all records for the Group to a new name. <b>Rename</b> is disabled when the name you are typing in the text box already exists in the CMDB.</p> <p><b>Merge</b> Changes all records for the Group to a different, already existing name in the CMDB. <b>Merge</b> is enabled when the name you are typing in the text box already exists in the CMDB.</p> <p><b>Move</b> Changes all records for the Group to a different, already existing name in the CMDB that you choose from the <b>New Group</b> drop-down menu.</p>
<b>Service</b>	Select a Service to edit, then click <b>Update</b> .

<b>Manage Service</b>	<p>Opens a dialog that enables you to <b>Delete</b>, <b>Rename</b> or <b>Merge</b> the Service.</p> <p><b>Delete</b> removes the Service from the CMDB database as well as all CMDB data and CIs associated with the Service.</p> <p><b>Rename</b> Changes all records for the Service to a new name. <b>Rename</b> is disabled when the name you are typing in the text box already exists in the CMDB.</p> <p><b>Merge</b> Changes all records for the Service to a different, already existing name in the CMDB. <b>Merge</b> is enabled when the name you are typing in the text box already exists in the CMDB.</p> <p><b>Move</b> Changes all records for the Service to a different, already existing name in the CMDB that you choose from the <b>New Service</b> drop-down menu.</p>
-----------------------	---

### CI List Table

This table lists all CIs associated with the selected Service. Each table row is a different CI. Select a CI to see its attributes in the drop-down menus at the right of the table. Use the **OSType**, **Region**, **SiteName**, **Criticality**, **City** and **Country** drop-down assign attributes, then click **Update**. To associate CIs with the Service, select one or more CIs from the **Available Components** table, then click **Add CI** (to associate the CI(s) with the selected Service.) or **Add CI To...** (to create a new Service and associate the CI(s) with it).

<b>CIType</b>	The type of CI. For example, server or application.
<b>CIName</b>	A CI name.
<b>Criticality</b>	<p>The importance level of the CI in your organization. Values range from <b>A</b> to <b>E</b>, where <b>A</b> is the highest Criticality and <b>E</b> is the lowest Criticality (with equally spaced intermediate values). This value is used to calculate the Alert Impact (maximum Alert Severity multiplied by the maximum Criticality equals Alert Impact).</p> <p>Criticality values are listed in the <b>Component Views - CI Service Table</b> display. Criticality values are also shown in heatmaps and tables.</p>
<b>Region</b>	The name of the Region for the CI.
<b>Environment</b>	The name of the Environment for the CI.
<b>SiteName</b>	The name of the Site for the CI.
<b>OSType</b>	The operating system on the CI.
<b>City</b>	The name of the City for the CI.
<b>Country</b>	The name of the Country for the CI.
<b>Update Criticality like selected CI</b>	Updates the Criticality attribute assigned to all CIs in the <b>CI List</b> table to match the selected CI level.
<b>Environ</b>	Select or type the Environment for the CI selected in the <b>CI List Table</b> , or the CI selected in the <b>Available Components</b> and added into the <b>CI List Table</b> .
<b>Region</b>	Select or type the region for the CI selected in the <b>CI List Table</b> , or the CI selected in the <b>Available Components</b> and added into the <b>CI List Table</b> .
<b>SiteName</b>	Select or type the site name for the CI selected in the <b>CI List Table</b> , or the CI selected in the <b>Available Components</b> and added into the <b>CI List Table</b> .
<b>Criticality</b>	<p>Specify the importance level of a Service or a CI for your organization. Select a Service or a CI and set the Criticality value from <b>A</b> to <b>E</b>, where <b>A</b> is the highest Criticality and <b>E</b> is the lowest Criticality (with equally spaced intermediate values). This value is used to calculate Alert Impact (maximum Alert Severity multiplied by the maximum Criticality equals Alert Impact).</p> <p>Criticality values are listed in the <b>Component Views - CI Service Table</b> display. Criticality values are also shown in heatmaps and tables.</p>

<b>Country</b>	Select or type the country for the CI selected in the <b>CI List Table</b> , or the CI selected in the <b>Available Components</b> and added into the <b>CI List Table</b> .
<b>OSType</b>	Select or type the operating system for the CI selected in the <b>CI List Table</b> , or the CI selected in the <b>Available Components</b> and added into the <b>CI List Table</b> .
<b>Update</b>	Updates the CI selected in the <b>CI List Table</b> with attributes selected from the drop-down menus (on the right).
<b>Delete</b>	Removes the selected CI from the CMDB database.

### Available Components Table

This table lists all available CIs in your RTView Enterprise system whether they are in the CMDB or not. Each row in the table is a different CI (for example, a server or a process). Select one or more CIs to associate with the currently selected Service, then click **Add CI** (to associate the CI(s) with the selected Service.) or **Add CI To...** (to create a new Service and associate the CI(s) with it).

<b>Selected CI Type</b>	Select the type of CI to include in the <b>Available Components</b> table. All CIs of this type are listed. A CI can be associated with multiple Services.
<b>CIName Filter</b>	Enter a string to filter the list of available components.
<b>Regex</b>	Check to enable Regex filtering.
<b>Add CI</b>	<p>Associates the CI selected in the <b>Available Components</b> table with the selected Service, and applies the attributes selected from the drop-down menus (on the right) to the CI.</p> <p>To associate a CI with the currently selected Service, select a CI from the <b>Available Components</b> table, use the drop-down menus on the right (<b>Environ</b>, <b>Region</b>, <b>SiteName</b>, etc.) to modify attributes for the CI, click <b>Add CI</b> and then click <b>Update</b>. The CI appears in the <b>CI List Table</b>.</p>
<b>Add CI To...</b>	<p>Creates a new Service and associates the selected CI with it.</p> <p>To create a new Service and associate a CI with it, select a CI from the <b>Available Components</b> table, use the drop-down menus on the right (<b>Environ</b>, <b>Region</b>, <b>SiteName</b>, etc.) to modify attributes for the CI, click <b>Add CI To...</b>, enter the name of the new Service, then click <b>Update</b>. The new Service is added to the list of Services and the CI appears in the <b>CI List Table</b>.</p>

## CMDB Administration - HTML

Manage and modify your Service Data Model (CMDB). You can add, delete, rename and merge CMDB hierarchical elements (Owners, Areas, Groups and Services). You can also associate a component ID (or CI) with Services and modify CI properties (such as the **Criticality** value). This display requires administrator privileges.

The screenshot shows the CMDB Administration interface. At the top, it displays the title "CMDB Administration" and the date/time "07-May-2019 15:56" with a "DATA" indicator. Below this, there are four configuration rows, each with a dropdown menu and a "Manage" button:

- Owner: APPLICATIONS (Manage Owner)
- Area: DEMO-APPS (Manage Area)
- Group: CLASSIC (Manage Group)
- Service: MYDEMO (Manage Service)

Below the configuration section is a table titled "CIs in Service" with an "Environment" dropdown set to "- All -". The table has the following columns: CI Type, CI Name, Environment, Criticality, Region, SiteName, City, Country, and OSType. The data rows are:

CI Type	CI Name	Environment	Criticality	Region	SiteName	City	Country	OSType
JVM	localhost.MYDEMO-DATA-1	PRODUCTION	C	AMER	Headquarters	Corte Merosa	California	Windows
JVM	localhost.MYDEMO-DB-1	PRODUCTION	C	AMER	Headquarters	Corte Merosa	California	Windows
JVM	localhost.MYDEMO-DISP-1	PRODUCTION	C	AMER	Headquarters	Corte Merosa	California	Windows

At the bottom of the table, there are three buttons: "Delete", "Edit...", and "Add CIs..."

Select a **Service** using the drop-down menus. The **CIs in Service** table lists CIs that your administrator has already associated with the selected **Service**.

Search by clicking the right side of a column heading > **Filter** to open the **Search, Sort and Column Selection** dialog.

The screenshot shows the "Search, Sort and Column Selection" dialog. It has a "Filter" dropdown menu on the left. The main area contains the following options:

- Show items with value that: Contains
- Sort Ascending
- Sort Descending (highlighted)
- Columns

At the bottom, there are "Filter" and "Clear" buttons.

**Export to Excel** by right-clicking a column heading.

By default, the Owner named **Infrastructure** is created. **Infrastructure** organizes CIs collected from Data Servers that are configured under RTView Enterprise by technology (for example, VMWARE-VM, TOMCAT-APP and EMS-QUEUE). This default organization can be disabled.

See instructions to:

- "Add CIs to the CMDB"
- "Delete a CI"
- "Edit CI Properties"
- "**Create, Delete, Rename or Merge Owners, Areas, Groups or Services**"
- "**Move a Service, Group or Area**"

For more details about the CMDB, see the *RTView Enterprise Configuration Guide*.

## Add CIs to the CMDB

You add CIs to the CMDB by associating them with a Service. CIs can be associated with more than one Service.

### To add CIs to the CMDB:

1. Click **Add CIs** to open the **Find CIs to add** table which contains all CIs that are available in your RTView Enterprise system (regardless of whether they are already in the CMDB).
2. Select one or more CIs in the **Find CIs to add** table and **Set CI properties** (including **Criticality**) using the drop-down menus. You can filter the list using the **CI Type** drop-down or by entering a search string in **CI Name Filter**.
3. Choose the Service you want to add the CI(s) to. You can:
  - add the CI(s) to an existing Service by selecting it from the **Service** drop-down (at the top of the display) and clicking **Add to Selected Service**.
  - add the CI(s) to a new Service you create by entering a **New Owner Name, New Area Name, New Group Name** and a **New Service Name**, and clicking **Add to New Service**.

The CI(s) are now listed in the **CIs in Service** table.

It is not necessary to restart the Configuration Server after making changes to the Service Data Model using the **CMDB Admin** display.

## Delete a CI

Select one or more CIs in the **CIs in Service** table, then click **Delete**. The CI is removed from the CMDB database and displays. Your changes are immediately visible in the drop-down menus and displays.

There is no option to undo a deletion from the CMDB. To restore a deletion you must find the CI again and re-associate it with the Service.

And when you delete all CIs from the list, the Service is also removed from the CMDB. A given Service can only exist if it contains one or more CIs. If the Service no longer exists as a result of removing the last of its CIs, you must also recreate the Service (by typing the names of the Owner, Area, Group, and Service).

## Edit CI Properties

Select one or more CIs in the **CIs in Service** table, then click . Use the drop-down menus to modify settings, then click . Your changes are immediately visible in the drop-down menus and displays.

Criticality is the importance level of a CI to your organization. Criticality values range from **A** to **E**, where **A** is the highest Criticality and **E** is the lowest Criticality (with equally spaced intermediate values). This value is used to calculate **Alert Impact** (the maximum Alert Severity multiplied by the maximum Criticality equals Alert Impact).

Criticality values are listed in the **Component Views - CI Service Table** display. Criticality values are also shown in heatmaps and tables.

## Create, Delete, Rename or Merge Owners, Areas, Groups or Services

You can create, delete and rename the **Owners, Areas, Groups** and **Services**. To illustrate, we use Owner as an example.

Select the Owner you want to modify and click the  (**Manage** cogwheel icon) next to it. You can:

- Delete the Owner by clicking .

This removes the Owner and all CI associations from the CMDB.
- Rename the Owner by entering a **New Owner Name** and clicking .

This changes the name of the Owner, creates a new Owner and retains all CI associations in the CMDB under the new Owner name.
- Create an Owner by clicking  to open the **Find CIs to add** table, then select one or more CIs in the **CIs in Service** table and enter the new Owner name in the **New Owner Name** field. Enter either an existing name or new name for the **Area, Group** and **Service** fields. Then click **Add to New Service** or **Add to Existing Service**.
- Merge all CIs under an Owner with another existing Owner by entering the existing target Owner in the **New Owner Name** field and clicking **Merge With Existing Owner**.

This changes the Owner name to that of the target Owner's name and moves all lower level CMDB associations (Services, Groups and Areas and associated CIs) go with it. For example, let's say **Owner A** is associated with one Area, that Area is associated with two different **Groups**, and both of those Groups are associated with two different Services. When you merge Owner A with Owner B, Owner B becomes the Owner of that one Area, the two Groups, the four Services and all the CIs associated with them.

This option is useful when, for example, an *existing* Owner is taking over for a retiring Owner.

## Move a Service, Group or Area

When you move a Service, Group or Area (Owners cannot be moved) you move it up one level in the CMDB and all lower level CMDB associations (Services, Groups and Areas and associated CIs) go with it.

This option is useful when, for example, it makes more organizational sense to have an Area under a different Owner, a Group under a different Area or a Service under a different Group. To illustrate, we use Group as an example.

Select the Group that you want to move to another Group and click the  (**Manage Group** cogwheel icon) next to it. You can:

- **Move all CIs to another Area** by selecting the target **Owner** and **Area** that you want to move the Group to and clicking .

This changes the name of the Owner and Area and retains all CI associations in the CMDB under the new Owner and Area.

- **Move all CIs to New Area** by typing the **New Owner Name** and **New Area Name** that you want to move the Group to and clicking .

This changes the name of the Owner and Area, creates a new Owner and Area and retains all CI associations in the CMDB under the new Owner and Area.

### Fields and Data

This display includes:

<b>Owner</b>	Select an Owner to filter by. The Owner selected populates the <b>Area</b> , <b>Group</b> and <b>Service</b> drop-down menus.
<b>Manage Owner</b>	<p>Opens a dialog that enables you to <b>Delete</b>, <b>Rename</b> or <b>Merge</b> the Owner.</p> <p><b>Delete</b> removes the Owner from the CMDB database as well as all CMDB data and CIs associated with the Owner.</p> <p><b>Rename</b> Changes all records for the Owner to a new name. <b>Rename</b> is disabled when the name you are typing in the text box already exists in the CMDB.</p> <p><b>Merge</b> Changes all records for the Owner to a different, already existing name in the CMDB. <b>Merge</b> is enabled when the name you are typing in the text box already exists in the CMDB.</p> <p><b>Note:</b> You cannot move Owners.</p>
<b>Area</b>	Select an Area to filter by. The Area selected populates the <b>Group</b> and <b>Service</b> drop-down menus.
<b>Manage Area</b>	<p>Opens a dialog that enables you to <b>Delete</b>, <b>Rename</b> or <b>Merge</b> the Area.</p> <p><b>Delete</b> removes the Area from the CMDB database as well as all CMDB data and CIs associated with the Area.</p> <p><b>Rename</b> Changes all records for the Area to a new name. <b>Rename</b> is disabled when the name you are typing in the text box already exists in the CMDB.</p> <p><b>Merge</b> Changes all records for the Area to a different, already existing name in the CMDB. <b>Merge</b> is enabled when the name you are typing in the text box already exists in the CMDB.</p> <p><b>Move</b> Changes all records for the Area to a different, already existing name in the CMDB that you choose from the <b>New Area</b> drop-down menu.</p>
<b>Group</b>	Select a Group to filter by. The Group selected populates the <b>Service</b> drop-down menu.

**Manage Group** Opens a dialog that enables you to **Delete**, **Rename** or **Merge** the Group.

**Delete** removes the Group from the CMDB database as well as all CMDB data and CIs associated with the Group.

**Rename** Changes all records for the Group to a new name. **Rename** is disabled when the name you are typing in the text box already exists in the CMDB.

**Merge** Changes all records for the Group to a different, already existing name in the CMDB. **Merge** is enabled when the name you are typing in the text box already exists in the CMDB.

**Move** Changes all records for the Group to a different, already existing name in the CMDB that you choose from the **New Group** drop-down menu.

**Service** Select a Service to edit, then click **Manage Service**.

**Manage Service** Opens a dialog that enables you to **Delete**, **Rename** or **Merge** the Service.

**Delete** removes the Service from the CMDB database as well as all CMDB data and CIs associated with the Service.

**Rename** Changes all records for the Service to a new name. **Rename** is disabled when the name you are typing in the text box already exists in the CMDB.

**Merge** Changes all records for the Service to a different, already existing name in the CMDB. **Merge** is enabled when the name you are typing in the text box already exists in the CMDB.

**Move** Changes all records for the Service to a different, already existing name in the CMDB that you choose from the **New Service** drop-down menu.

#### CI's in Service Table

This table lists all CIs associated with the selected Service. Each table row is a different CI. Select a CI and click **Edit** to revise its properties and **Delete** to remove the CI from the CMDB. Click **Add CI's...** to open the **Find CI's to add** table and select the CIs you want to add to/associate with the selected Service.

**CIType** The type of CI. For example, server or application.

**CIName** A unique identifier for the CI.

**Criticality** The importance level of the CI in your organization. Values range from **A** to **E**, where **A** is the highest Criticality and **E** is the lowest Criticality (with equally spaced intermediate values). This value is used to calculate the Alert Impact (maximum Alert Severity multiplied by the maximum Criticality equals Alert Impact).  
Criticality values are listed in the **Component Views - CI Service Table** display. Criticality values are also shown in heatmaps and tables.

**Region** The name of the Region for the CI.

**Environment** The name of the Environment for the CI.

**SiteName** The name of the Site for the CI.

**OSType** The operating system on the CI.

**City** The name of the City for the CI.

**Country** The name of the Country for the CI.

- Criticality** Specify the importance level of a Service or a CI for your organization. Select a Service or a CI and set the Criticality value from **A** to **E**, where **A** is the highest Criticality and **E** is the lowest Criticality (with equally spaced intermediate values). This value is used to calculate Alert Impact (maximum Alert Severity multiplied by the maximum Criticality equals Alert Impact).  
Criticality values are listed in the **Component Views - CI Service Table** display. Criticality values are also shown in heatmaps and tables.
- Country** Select or type the country for the CI selected in the **CI List Table**, or the CI selected in the **Available Components** and added into the **CI List Table**.
- OSType** Select or type the operating system for the CI selected in the **CI List Table**, or the CI selected in the **Available Components** and added into the **CI List Table**.
- Delete** Removes the selected CI from the CMDB database.

#### Find CIs to add Table

This table opens when you click **Add CIs...** This table lists all available CIs in your RTView Enterprise system whether they are in the CMDB or not. Each row in the table is a different CI (for example, a server or a process). Select one or more CIs to associate with the currently selected Service, then click **Add to Selected Service** (to associate the CI(s) with the selected Service), or **Add To New Service** (to create a new Service and associate the CI(s) with it) or **Add to Existing Service** (to select a different, existing Service). The **CIs in Service Table** is updated with the CI additions immediately.

## Architecture

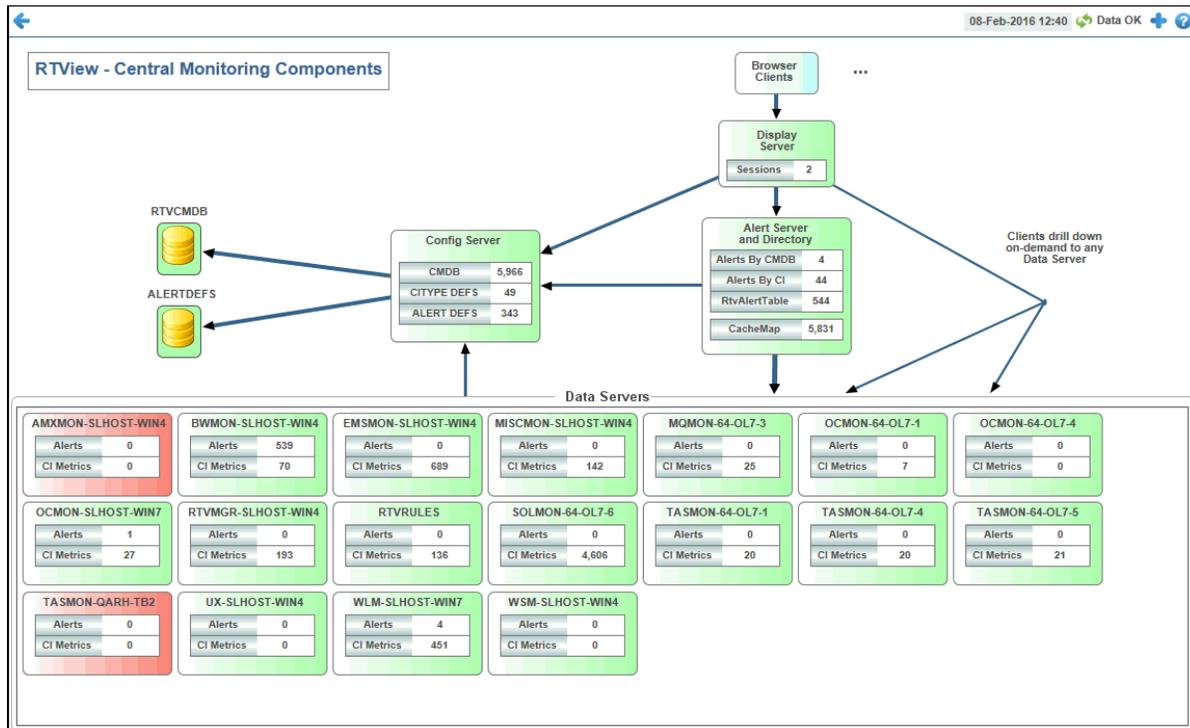
These displays provide a view of RTView Enterprise component connectivity, mapping between component types, and component level connection and performance information. The Architecture displays are provided with RTView Enterprise. Displays in this View are:

- **"System Overview"** Topology of the main RTView Enterprise components. Objects are color-coded to show component status.
- **"RTView Data Servers"**: Configuration and connection details for RTView Data Servers.
- **"Data Server Summary"**: Connection and query statistics for RTView Data Servers.
- **"RTView History Table Statistics"**: Performance of historical data being stored from caches with history.
- **"RTView Cache Tables"**: Configuration and alert details for RTView Cache Tables.
- **"RTView CI Stats Tables"**: Alert details for RTView Cache Tables by CI.
- **"RTView CI Type Defs"**: CI Type definitions, cache map and alert map by CI Type.
- **"RTView KM Defs"**: Key Metrics definitions for all CI Types.
- **"About"**: This display shows details about the RTView Enterprise version and data sources available to your system.

## System Overview

View the topology of the central RTView Enterprise monitoring components and their current connection state. Each object represents a component which are color-coded to indicate component status. Red indicates the component stopped running. Green indicates the component is running.

For details about the HTML version of this display, see ["System Overview - HTML"](#).



#### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

### Fields and Data

This display includes:

#### Config Server

The Configuration Server provides configurations to all central RTView Enterprise components.

- CMDB** The number of CIs in the CMDB.
- CITYPE DEFS** The current number of CITYPE definitions.
- ALERTDEFS** The current number of alert settings and override definitions.

#### Alert Server and Directory

The Alert and Directory Server centralizes access to all alerts sent by remote Data Servers, and maintains a directory table of CI locations. The CI location is the name of the source Data Server.

**Alerts By CMDB** The number of Services in the CMDB that currently have at least one associated alert.

**Alerts By CI** The number of CIs in the CMDB that currently have at least one associated alert.

**RtvAlertTable** The number of currently active alerts in the system.

**CacheMap** The number of entries currently in the directory table.

**Display Server**

The Display Server generates HTML displays for browser clients.

**Sessions** The current number of users connected to the Display Server.

**Browser Clients** The browser clients are represented in the topology as a single object. No data is shown for browser clients.

**Data Servers**

This panel in the topology shows all Data Servers.

**Alerts** The number of currently activated alerts for the Data Server.

**CI Metrics** The count of CI metrics that the remote Data Server is sending.

**RTView Data Servers**

View Data Server connection status and detailed client connection information. For details about the HTML version of this display, see "RTView Data Servers - HTML".

←
**RTView Data Server Tables**
23-Sep-2015 14:22 Data OK + ?

**Local Connections to DataServer**

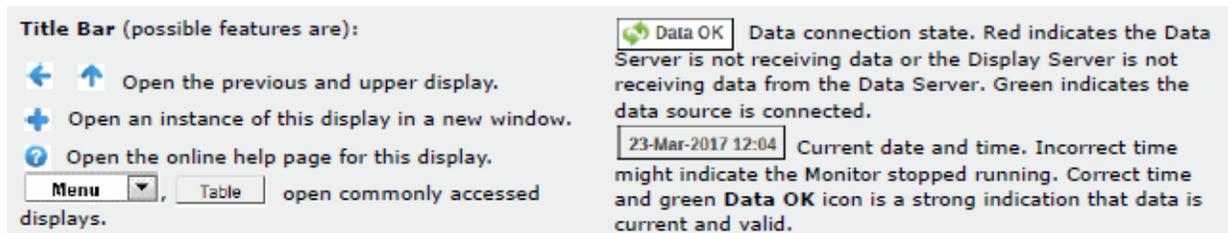
Name	Connected	Status	ConnectionString	Rcv Cnt	ReceiveT
ALERT_SERVER	<input checked="" type="checkbox"/>	OK	192.168.200.134:10028	5,111	9/23/15 14
AMXMON-SLHOST-WIN4	<input checked="" type="checkbox"/>	OK	192.168.200.134:6378	52	9/23/15 13
BW8MON-SLHOST-WIN4	<input checked="" type="checkbox"/>	OK	192.168.200.134:3378	5	9/23/15 13
CONFIG_SERVER	<input checked="" type="checkbox"/>	OK	192.168.200.134:10018	133	9/23/15 14
EMSMON-SLHOST-WIN4	<input checked="" type="checkbox"/>	OK	192.168.200.134:3178	5	9/23/15 13
MISCMON-SLHOST-WIN4	<input checked="" type="checkbox"/>	OK	192.168.200.134:3978	112	9/23/15 11
MQMON-64-OL7-3	<input checked="" type="checkbox"/>	OK	192.168.200.73:3478	3	9/23/15 11
OCMON-64-OL7-1	<input type="checkbox"/>	no conned	192.168.200.71:3381	0	12/31/69 16
OCMON-64-OL7-4	<input checked="" type="checkbox"/>	OK	192.168.200.74:3381	3	9/23/15 11
OCMON-SLHOST-WIN7	<input type="checkbox"/>	no conned	192.168.200.137:3381	0	12/31/69 16
RTVMGR-SLHOST-WIN4	<input checked="" type="checkbox"/>	OK	192.168.200.134:3078	33	9/23/15 13

**DataServer Manager**

NumberOfClients	ServingData	ConnectionRequestCount	ConnectionRequestFailedCount
5	<input checked="" type="checkbox"/>	5	

**DataServer Clients**

Client ID	Address	Host	Process Name	PID	Last Data Sent	Total Data Sent	Du
3	192.168.200.134	SLHOST-WIN4	dataserverd	55364@SLHOST-WIN4	5,258	93,963,548	0 04:
4	127.0.0.1	127.0.0.1	historiand	15868@SLHOST-WIN4	2,722	37,981,383	0 04:
1	192.168.200.134	SLHOST-WIN4	displayserver	27116@SLHOST-WIN4	190,604	616,260,790	0 04:
2	127.0.0.1	127.0.0.1	dataserverd	10564@SLHOST-WIN4	17,287	126,666,689	0 04:
5	192.168.200.134	SLHOST-WIN4	displayserver	55200@SLHOST-WIN4	5,258	73,499,814	0 04:



## Fields and Data

This display includes:

### Local Connections to Data Server

This table lists all Data Servers and detailed connection information. Select a Data Server to view further details (in the lower tables).

<b>Name</b>	The Data Server name.
<b>Connected</b>	When checked, the connection is currently connected.
<b>Status</b>	The Data Server connection status.
<b>Connection String</b>	The host name and port number for TCP connections, or the URL for servlet connections.
<b>Rcv Cnt</b>	The number of data updates received from that Data Server.
<b>ReceiveTime</b>	The time that data was last received.
<b>Config</b>	The RTView version running on the Data Server.

### Data Server Manager

This table shows connection information for the Data Server selected from the **Local Connections to Data Server** table.

<b>NumberOf Clients</b>	The number of clients currently connected to the Data Server.
<b>ServingData</b>	When checked, the Data Server is currently serving data.
<b>Connection Request Count</b>	The number of client requests to connect to the Data Server.
<b>Connection Request FailedCount</b>	The number of client requests to connect to the Data Server that were unable to connect.

### Data Server Clients

This table shows information for clients connected to the Data Server selected from the **Local Connections to Data Server** table.

- ClientID**            A unique string identifier for the client.
- Address**            The client IP address.
- Duration**           The client session length of time.
- Host**                The address of the client host.
- Last Data Sent**     The amount of data, in bytes, the Data Server last sent to the client.
- Total Data Sent**    The total amount of data, in bytes, the Data Server has sent to the client.

## Data Server Summary

View Data Server connection status, cache table sizes and database query metrics. Use the available drop-down menus or right-click to filter data shown in the display. For details about the HTML version of this display, see ["Data Server Summary - HTML"](#)

**RTView Data Server - Summary**    23-Sep-2015 14:20    Data OK

Data Server: <Default>

Connection Status						RTView Cache Tables	
Connected	Status	Connection String	Receive Count	Receive Time	Config	CacheTable	Rows
						RtvAlertTableLocal	19,905
						RtvMxCacheDefsWithCo	1,429
						RtvTabTreeCache	488
						RtvMxCacheDefsRaw	234
						RtvMxCacheDefs	169
						RtvCmdbServiceTable	56
						RtvMxCacheInfoByServ	56
						RtvDataServerConnecti	20
						RtvCmdbGroupTable_I	17
						RtvCmdbAreaTable_loc	8
						RtvCmdbOwnerTable_I	2
						JmxStatsTotals	1
						RtvAlertMapByCI	0
						RtvAlertSourceStats	0
						RtvAlertStatsByCI	0
						RtvAlertStatsByCIAndA	0

Database Queries (running on selected Data Server)								
Database	Conn	Count	Active	ExecTime	Rows	RunTime	Status	
ALERTDEFS	<input checked="" type="checkbox"/>	0	<input type="checkbox"/>	NaN	0			
PROPDB	<input checked="" type="checkbox"/>	0	<input type="checkbox"/>	NaN	0			
RTVCMDB	<input checked="" type="checkbox"/>	0	<input type="checkbox"/>	NaN	0			
RTVCONFIG	<input checked="" type="checkbox"/>	0	<input type="checkbox"/>	NaN	0			
RTVHISTORY	<input checked="" type="checkbox"/>	0	<input type="checkbox"/>	NaN	0			

**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

Data OK    Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04    Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Fields and Data**

This display includes:

**Data Server** Select a Data Server from the drop-down menu to view details for in the display.

**Connection Status**

This table shows connection details for the selected Data Server.

<b>Connected</b>	When checked, the Data Server is currently connected.
<b>Status</b>	The Data Server connection status.
<b>Connection String</b>	The host name and port number for TCP connections, or the URL for servlet connections.
<b>Rcv Cnt</b>	The number of data updates received from that Data Server.
<b>ReceiveTime</b>	The time that data was last received.
<b>Config</b>	The RTView version running on the Data Server.

**Alert Table View** Select to view or manage current alerts for the selected Data Server in the **RTView Alerts Table** display.

**Alert Admin** Select to view or manage alert thresholds for the selected Data Server in the **Alert Administration** display.

**History Tables** Select to view database table statistics for each cache for the selected Data Server in the ["RTView History Table Statistics"](#) display.

**RTView Cache Tables**

This table lists Cache Tables and their size, in number of rows, for the selected Data Server. Select a Cache Table to view details in the **RTView Cache Tables** display.

Use this data for debugging. This display is typically used for troubleshooting with SL Technical Support.

<b>CacheTable</b>	The name of the Cache Table.
<b>Rows</b>	The current number of rows in the Cache Table.

**Database Queries**

This table lists the databases and query details for the selected Data Server. Each table row describes a different query.

<b>Database</b>	The name of the database.
<b>Conn</b>	When checked, the database is currently connected.
<b>Count</b>	The number of query requests from current Data Server.
<b>Active</b>	When checked, the query is currently running.
<b>ExecTime</b>	The amount of time, in milliseconds, to execute the query.
<b>Rows</b>	The number of rows the query created.
<b>RunTime</b>	The time the query was executed.
<b>Status</b>	The latest result status of the query.
<b>Query</b>	The query that was executed.

## RTView History Table Statistics

This display opens when you click **History Tables** from the **Architecture - "Data Server Summary"** display. View information about the performance of historical data being stored from caches with history. Use this display to verify your tables are growing as expected by:

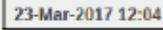
- seeing how many rows are in the database table (**Row Count**).
- seeing how many rows are added at each update period (**Delta**).
- verifying that the range of the data stored in the table is consistent with defined compaction rules and that behavior is as expected. To do this, compare the time of **First Entry** and **Last Entry** and verify the dates match the defined compaction interval (for example, **2 weeks** by default). For this verification, you must first confirm the historian has been operating for at least the defined compaction time interval, otherwise the range of data will be shorter.

Cache Name / DB Table Name		Row Count	Delta	Distinct	First Entry	Last Entry	Current
EmsAdmStats EMS_ADMSTATS		8,337	1	0	18-Sep-2014 01:15:00	10-Dec-2014 07:16:00	🟢
EmsDurables EMS_DURABLES		395,397	53	53	18-Sep-2014 01:15:00	10-Dec-2014 07:15:00	🟢
EmsQueueTotalsByServer EMS_QUEUETOTALS		181,236	181,236	24	18-Sep-2014 01:15:00	10-Dec-2014 05:25:00	🔴
EmsQueues EMS_QUEUES		3,958,595	296	665	15-Sep-2014 02:00:00	10-Dec-2014 07:16:00	🟢
EmsRouteCountsByServer EMS_ROUTECOUNTS		124,790	15	16	18-Sep-2014 01:15:00	10-Dec-2014 07:16:00	🟢
EmsRoutes EMS_ROUTES		208,014	25	26	18-Sep-2014 01:15:00	10-Dec-2014 07:16:00	🟢
EmsServerInfo EMS_SERVERINFO		199,807	24	25	18-Sep-2014 01:15:00	10-Dec-2014 07:16:00	🟢
EmsTopicTotalsByServer EMS_TOPICTOTALS		183,187	22	23	18-Sep-2014 01:15:00	10-Dec-2014 07:16:00	🟢
EmsTopics EMS_TOPICS		7,132,737	621	37,369	18-Sep-2014 01:15:00	10-Dec-2014 07:16:00	🟢
JvmMemory JVM_MEMORY		20,737	0	13	27-Nov-2014 00:15:00	10-Dec-2014 07:13:00	🟢
JvmOperatingSystem JVM_OPERATINGSYSTEM		20,214	17	13	27-Nov-2014 00:15:00	10-Dec-2014 07:21:00	🟢

### Title Bar (possible features are):

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu ▾, Table open commonly accessed displays.

 Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

### Fields and Data

This display includes:

- Cache Name / DB Table Name** The name of the cache and the name of the database table. Mouse-over to see the **Index** columns for the cache.
- Row Count** The number of rows in the table.
- Delta** The number of rows added since the last update.
- Distinct** The number of distinct indexes in the table.

- First Entry** The time stamp of the oldest entry written to the table.
- Last Entry** The time stamp of the most recent entry written to the table.
- Current** The current writing state of the table.
  - (Time ≥ 10m) The writing latency is equal to or greater than ten minutes.
  - (Time ≥ 4m and < 10m) The writing latency is equal to or greater than four minutes and less than ten minutes.
  - (Time < 4m) The writing latency is less than four minutes.

## RTView Cache Tables

For details about the HTML version of this display, see ["Cache Table - HTML"](#)

View Data Server Cache table sizes and contents. Select a cache table in the upper table and view its contents in the lower table. Use the available drop-down menus or right-click to filter data shown in the display.

**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Fields and Data**

This display includes:

- Data Server** Select a Data Server from the drop-down menu to view details for in the display.
- Max Rows** Enter the maximum number of rows to include in the lower table, then click Enter.
- History Tables** Select to include all defined history tables in the **RTView Cache Tables** list.

**RTView Cache Tables**

This table lists cache tables for the selected Data Server. Select a cache table to view details in the lower table.

<b>CacheTable</b>	The name of the cache table.
<b>TableType</b>	The type of cache table.
<b>current</b>	This table is a current table which shows the current values for each index.
<b>current_condensed</b>	This table is a current table with primary compaction configured.
<b>history</b>	This table is a history table.
<b>history_condensed</b>	This table is a history table with primary compaction configured.
<b>history_combo</b>	This table is a history table with primary compaction configured, and which is also configured to store rows of recent raw data followed by rows of older condensed data.
<b>Rows</b>	The number of rows currently in the table.
<b>Columns</b>	The number of columns currently in the table.
<b>Memory</b>	The amount of space, in bytes, used by the table.

**(Lower Table)**

This table shows the contents of the selected cache table. Available columns vary by cache. For example, a JVM cache table might provide **BootClassPath** and **InputArgument** columns, and a Tomcat cache might provide **RateAccess** and **cacheMaxSize** columns.

<b>Rows</b>	The number of rows currently in the table.
-------------	--

## RTView CI Stats Tables

View details for components that currently have an active warning or alarm alert. For details about the HTML version of this display, see "[CI Stats - HTML](#)"

Alert Stats By CI				
time_stamp	CItYPE	CInAME	MaxSeverity	AlertCount
09/25/15 11:13:33	BW-ENGINE	slhpux11(simon);domainsimon.BWApp-7.Procs	1	1
09/25/15 11:13:33	BW-ENGINE	slapm(slapm);domainslapm.BWApp-5.Procs	2	1
09/25/15 11:13:33	BW-ENGINE	slvmr2(slapm);domainslapm.BWEngine.Process Archive	2	1
09/25/15 11:13:33	BW-ENGINE	slie4-64(simon);domainsimon.BWApp-4.Procs	2	1
09/25/15 11:13:33	BW-ENGINE	slapm(slapm);domainslapm.BW Engine.Process Archive	2	1
09/25/15 11:13:33	BW-ENGINE	slapm(slapm);domainslapm.BWApp Spate.Procs-1	2	1
09/25/15 11:13:33	BW-ENGINE	slie4-64(simon);domainsimon.BWApp-5.Procs	1	1
09/25/15 11:13:33	BW-ENGINE	slie4-64(simon);domainsimon.BWApp-10.Procs	1	1

Cache Map By CItYPE			
time_stamp	CItYPE	CACHENAME	Source
09/25/15 09:33:53	ACW	AwsEc2InstanceStats	MISCMON-DATA-1
09/25/15 01:32:30	BW-ENGINE	BwEngines	Z-SIMDATA-1
09/25/15 06:02:09	BW-ENGINE	BwEngines	BWMON-SLDEMOS
09/25/15 01:32:30	BW-PROCESS	BwProcesses	Z-SIMDATA-1
09/25/15 06:02:09	BW-PROCESS	BwProcesses	BWMON-SLDEMOS
09/25/15 01:32:30	BW-SERVER	BwServers	Z-SIMDATA-1
09/25/15 06:02:09	BW-SERVER	BwServers	BWMON-SLDEMOS

Cache Map By CI Count: 2475

### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

### Fields and Data

This display includes:

#### Alert Stats By CI

This table provides summary alert details for all CIs that currently have active warning or alarm alerts.

<b>time_stamp</b>	The date and time this table row of data was last updated. Format: <b>MM/DD/YY HH:MM:SS</b> <b>&lt;month&gt;/ &lt;day&gt;/&lt;year&gt; &lt;hours&gt;:&lt;minutes&gt;:&lt;seconds&gt;</b>
<b>CItYPE</b>	The component type.
<b>CInAME</b>	The name of the component.
<b>MaxSeverity</b>	The most critical alert state of all current alerts for this component.
<b>AlertCount</b>	The number of current warning and alarm alerts for this component.

#### Cache Map By CItYPE

This table provides mapping of all component types to caches.

<b>time_stamp</b>	The date and time this table row of data was last updated. Format: <b>MM/DD/YY HH:MM:SS</b> <b>&lt;month&gt;/ &lt;day&gt;/&lt;year&gt; &lt;hours&gt;:&lt;minutes&gt;:&lt;seconds&gt;</b>
<b>CItYPE</b>	The component type.
<b>CACHENAME</b>	The name of the cache associated with the component type.

**Source** The name of the Data Server alert sending data for that component type.

**Cache Map By CI**

This table provides the location of all CIs.

**CI Type Filter:** Select the CI Type to filter by in this table, or select **All CI Types**.

**Count** The number of CIs currently in this table.

**time\_stamp** The date and time this table row of data was last updated.  
 Format:  
**MM/DD/YY HH:MM:SS**  
**<month>/ <day>/<year> <hours>:<minutes>:<seconds>**

**CIType** The component type.

**CIName** The name of the component.

**DataServerName** The name of the Data Server which sent this CI.

**Expired** When checked, data has not been received within the time specified in the solution package that is hosting the data. If the solution package is configured to delete expired data, this row will be deleted if no data is received within the time specified for deletion. See the documentation for the solution package that is hosting the data for information on how to configure expiration and deletion times.

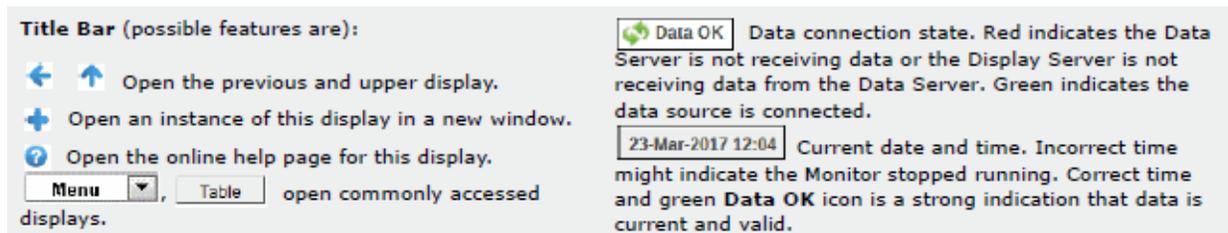
**RTView CI Type Defs**

This display provides component type definitions and shows the mapping of component types to caches as well as component types to alerts. For details about the HTML version of this display, see ["CI Type Definitions - HTML"](#)

RTView CI Type Definitions						
CITYPE	INDEXMAP	INDEXNAMES	RTVDISPLAY	CIVARMAP	DEFAULTQUALITY	OW
ACW	1	Dimension	acw_instance_summary	\$awsEc2InstanceId		1 Infrastr
AMX-HOST	1	AMX Host	amx_host_summary	\$amxHost		1 Infrastr
AMX-NODE	1,2	AMX Host;Node	amx_node_summary	\$amxHost;\$amxNode		1 Infrastr
AMX-SERVICE	1,2	Application;Service	amx_service_summary	\$amxApplication;\$amxService		1 Infrastr
AMX-SERVICENODE	1,2,3,4	AMX Host;Node;Application;Ser...	amx_servicenode_summary	\$amxHost;\$amxNode;\$amxApp...		1 Infrastr
BW6-APP	1,2,3	Domain;AppSpace;Application...	bw6_app_summary	\$bw6domain;\$bw6appspace;\$b...		1 Infrastr
BW6-APPNODE	1,2,3	Domain;AppSpace;AppNode	bw6_appnode_summary	\$bw6domain;\$bw6appspace;\$b...		1 Infrastr
BW6-PROCESS	1,2,3,4,5	Domain;AppSpace;AppNode;Ap...	bw6_process_summary	\$bw6domain;\$bw6appspace;\$b...		1 Infrastr
BW-ENGINE	1,2	AgentName;MicroAgentName	bw_engine_summary	\$bwserver;\$bwengine		1 Infrastr

Cache Map By Citype		Alert Map By Citype	
CITYPE	CACHENAME	CITYPE	ALERTNAME
ACW	AwsEc2InstanceStats	ACW	AcwInstanceCpuHigh
AMX-NODE	AmxNodes	ACW	AcwInstanceDiskReadBytesHigh
AMX-SERVICE	AmxServiceTotals	ACW	AcwInstanceDiskReadOpsHigh
AMX-SERVICEN...	AmxServices	ACW	AcwInstanceDiskWriteBytesHigh
BW6-APP	Bw6Apps	ACW	AcwInstanceDiskWriteOpsHigh
BW6-APP	Bw6ProcessTotalsByApp	ACW	AcwInstanceNetworkReadBytesHigh
BW6-APPNODE	Bw6AppNodes	ACW	AcwInstanceNetworkWriteBytesHigh
BW6-PROCESS	Bw6Processes	AMX-SERVICE	AmxServiceNodeHitRateHigh
BW-ENGINE	BwEngines	AMX-SERVICE	AmxServiceNodeResponseTimeHigh
BW-PROCESS	BwProcesses	AMX-SERVICE	AmxServiceNodeFaultRateHigh
BW-SERVER	BwServers	AMX-SERVICE	AmxServiceHitRateHigh
DB2	Db2DbSummary	AMX-SERVICE	AmxServiceResponseTimeHigh
DB2	Db2ResponseTime	AMX-SERVICE	AmxServiceFaultRateHigh
EM-SERVICE	RtvCmdServiceTable_Local	AMX-SERVICE	AmxServiceNodeNotRunning
EM-SERVICE	RtvCmdServiceStats_Local	BW6-APP	Bw6AppProcessCreatedRateHigh



## Fields and Data

This display includes:

### CI Type Definitions

This table provides definitions for all CI Types.

<b>CITYPE</b>	The component type.
<b>INDEXMAP</b>	Number of indexes and the order in which they are used to form the CI Name.
<b>INDEXNAMES</b>	Semicolon-separated list of the index columns.
<b>RTVDISPLAY</b>	The name of the RTView display to drill-down to from the <b>Alerts Table</b> to see summary data for this CI Type. This is the target of the <b>Go To CI</b> button in the <b>Alerts Table</b> and in the <b>Service Summary</b> display.
<b>CIVARMAP</b>	The names of substitutions that must be set to drill-down to the display.
<b>DEFAULTQUALITY</b>	A flag indicating whether the lack of data is considered an error condition or not.
<b>OWNER</b>	The Owner the CITYPE is associated with, when the CMDB is populated automatically from CIs of this type.
<b>AREA</b>	The Area the CITYPE is associated with.
<b>SERVICEGROUP</b>	The SERVICEGROUP the CITYPE is associated with, when the CMDB is populated automatically from CIs of this type.

### Cache Map By CITYPE

This table provides mapping of component types to caches for all component types.

<b>CITYPE</b>	The type of CI.
<b>CACHENAME</b>	The name of the cache associated with the component type.

### Alert Map By CITYPE

This table provides mapping of component types to alerts.

<b>CITYPE</b>	The type of CI.
<b>ALERTNAME</b>	The name of the alert.

## RTView KM Defs

This display shows the Key Metrics definitions for all CI Types. For details, see [“Available KM Metrics and Alerts”](#).

RTView Key Metrics Definitions				
CITYPE	CACHENAME	SELECTOR	METRICNAME	AlertName
ACW	AwsEc2InstanceStats	Instance CPU Usage	CPUUtilization	AowInstanceCpuHigh
AMX-SERVICE	AmxServiceTotals	Service Hits/Min	Hits Per Minute	AmxServiceHitRateHigh
AMX-SERVICE	AmxServiceTotals	Service Response Time	Avg. Response Time	AmxServiceResponseTir
AMX-SERVICENODE	AmxServices	Node Hits/Min	Hits Per Minute	AmxServiceNodeHitRate
AMX-SERVICENODE	AmxServices	Node Response Time	Avg. Response Time	AmxServiceNodeRespor
BW6-APP	Bw6ProcessTotalsByApp	App Created / sec	RateCreated	Bw6AppProcessCreatedf
BW6-APP	Bw6ProcessTotalsByApp	App Exec Time / sec	RateTotal Execution	Bw6AppProcessExecutio
BW6-APPNODE	Bw6AppNodes	CPU Used %	Used CPU Percentage	Bw6AppNodeCpuUsedH
BW6-APPNODE	Bw6AppNodes	Mem Used %	Used Memory Percentage	Bw6AppNodeMemUsedf
BW6-PROCESS	Bw6Processes	Process Created / sec	RateCreated	Bw6ProcessCreatedRate
BW6-PROCESS	Bw6Processes	Process Exec Time / sec	RateTotal Execution	Bw6ProcessExecutionTi
BW-ENGINE	BwEngines	CPU Used %	CPU %	BwEngineCpuUsedHigh
BW-ENGINE	BwEngines	Memory Used %	PercentUsed	BwEngineMemUsedHigh
BW-PROCESS	BwProcesses	Process Exec Time / sec	RateTotalExecution	BwProcessExecutionTim
BW-SERVER	BwServers	CPU Used %	CPU Usage %	BwServerCpuUsedHigh
DB2	Db2ResponseTime	Response Time	ResponseTimeMilliSec	Db2ResponseTimeHigh
EM-SERVICE	RtvCmdServiceStats_local	Alert Impact	AlertImpact	RtvEmServiceAlertImpa
EMS-QUEUE	EmsQueues	Pending Msgs	pendingMessageCount	EmsQueuesPendingMsg
EMS-QUEUE	EmsQueues	In Msgs / sec	inboundMessageRate	EmsQueuesInMsgRateH
EMS-QUEUE	EmsQueues	Out Msgs / sec	outboundMessageRate	EmsQueuesOutMsgRate
EMS-QUEUE	EmsQueues	Consumers	consumerCount	EmsQueuesConsumerC

### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu , open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

### Fields and Data

This display includes:

#### RTView Key Metrics Definitions

This table provides Key Metrics definitions for all CI Types.

<b>CITYpe</b>	The component type.
<b>CACHENAME</b>	The name of the cache that contains the Key Metric.
<b>SELECTOR</b>	The name used for this Key Metric in the <b>Metric Name</b> column of Key Metric displays.
<b>METRICNAME</b>	The name of the <b>cache</b> column that contains this Key Metric.
<b>ALERTNAME</b>	The name of the alert associated with this Key Metric. When blank, the Key Metric is not configured for inclusion in Key Metric displays.

**CalcMode**

The calculation used for the **Threshold %** value. The base value is calculated as the percent of the Key Metric value between **0** and the **ALARMLEVEL** of the associated alert. If the **CalcMode** is blank, this value is used. If the **CalcMode** is:

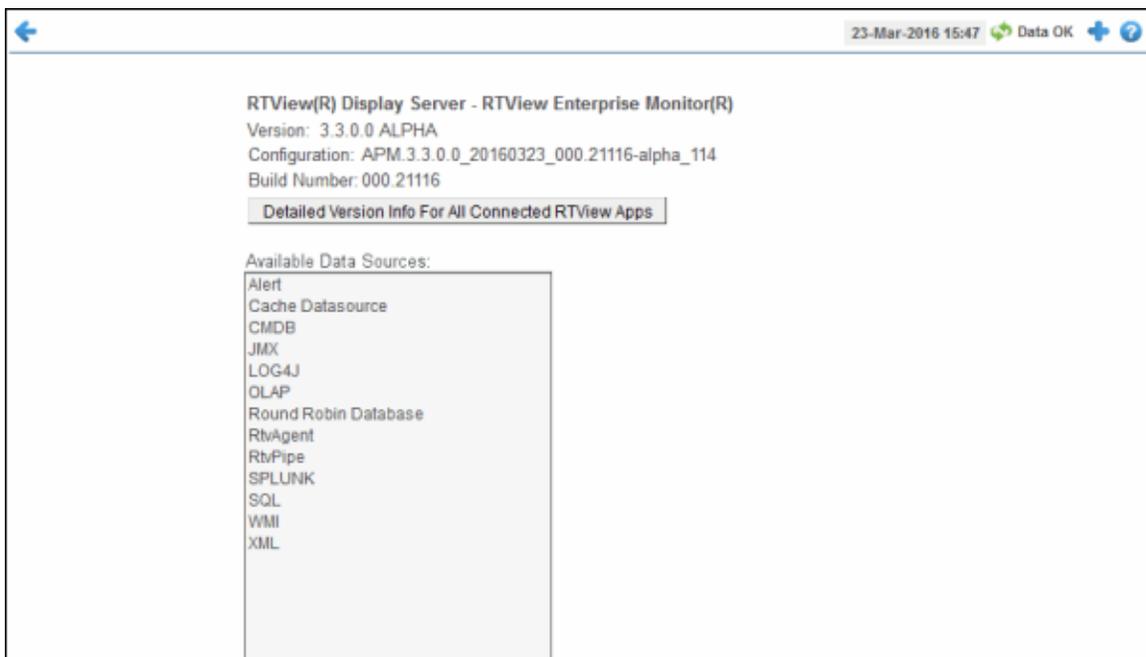
- **exp**, the value is adjusted so that lower values are diminished. Typically, this is used for memory metrics.
- **inverse**, the value is calculated in reverse of the standard thresholds. This is used when the associated alert is a low threshold alert.
- **invpct**, the value is calculated in reverse of the standard threshold and is assumed to be percent and therefore a value between **0** and **100**. This is used when the associated alert is a low threshold alert against a percent.
- **log**, a logarithmic algorithm is applied.

**Level**

The Key Metric level. Level **0** KMs are always displayed. Level **1** KMs are displayed if **Show More Metrics** is selected.

**About**

Get RTView Enterprise version and configuration information including a list of all available data sources.



## Property Views

These displays show how your properties are configured and the values for all connected RTView processes. The displays are located under the **ADMIN** tab. Displays in this View are:

- ["Properties Configuration"](#): Table of properties configuration settings, per connection.
- ["System Properties"](#): Table of system properties for RTView processes, per connection.
- ["Applied Properties"](#): Table of all properties that were applied to RTView processes, per connection.
- ["All Properties"](#): Table of all properties that were read from the properties files and database regardless of whether or not the RTView process uses them.
- ["Properties Descriptions"](#): Table of all properties that are supported by RTView processes, per connection.

### Properties Configuration

This display shows properties configuration information. The **Last Property Read Time** shows the last time that properties were read for the RTView process specified by the selected **Connection**.

Source: <input type="text" value="Data Server"/>	Connection: <input type="text" value="ALERT_SERVER"/>
Last Property Read Time: Mar 23, 2016 10:53:12 AM	

Select the **Source** of the connection to the RTView process for which you want to see property information. Options are:

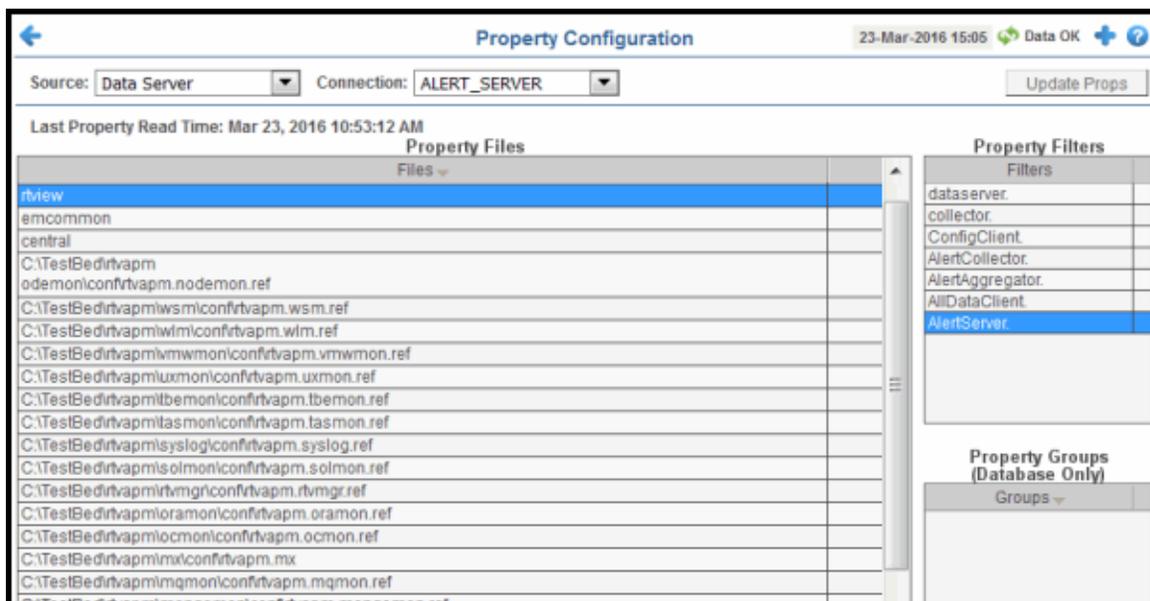
- **Data Server:** If the RTView process is a Data Server and the Thin Client has a defined Data Server connection for it, choose this option and select the name of the Data Server in the Connection field.
- **Local JMX Connection:** Select this option if the Thin Client has a defined JMX Connection to the RTView process.
- **RTVMGR JMX Connection:** Select this option if the RTView has a defined JMX Connection to the RTView process.

Select the **Connection** to the RTView process for which you want to see property information. Options available depend on your setup. For example, **RTView** is only visible when the **Source** is **RTView JMX Connection** and you have multiple RTView s. You can then select an RTView that has a defined JMX Connection to the RTView process for which you want to see property information.

The **Property Files** table shows all of the properties files that were read by the RTView process specified by the selected **Connection** in the order they were read. The **Property Filters** table shows all filters that are applied to the properties. **Property Groups** shows all property groups that are applied to the properties. **Property Groups** are only used when reading properties from a database.

Click **Update Props** to have the RTView process specified by the selected **Connection** re-read all properties files and database properties. Note that most non-connection properties do NOT support updates. See the ["Properties Descriptions"](#) display to find out if a specific property supports updates.

Right-click/**Export** to create a PDF image of the display. Click Sort  to order column data.



#### Title Bar (possible features are):

-   Open the previous and upper display.
-  Open an instance of this display in a new window.
-  Open the online help page for this display.
- open commonly accessed displays.

-  **Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Note:** The "Up" Arrow () opens the most recently viewed display under "Multi Area Service Views". For example, if the last viewed display under **Multi Area Service Views** was **Services CI Type Summary**, then clicking  opens the "Services CI Type Summary" display.

#### Filter By:

- Source:** Select the **Source** of the connection to the RTView process for which you want to see property information.
- Connection:** Select the **Connection** to the RTView process for which you want to see property information.

#### Fields and Data

This display includes:

- Update Props** Click to have the RTView process specified by the selected Connection re-read all properties files and database properties. Note that most non-connection properties do NOT support updates. Use the "Properties Descriptions" display to see if a specific property supports updates.
- Last Property Read Time** The last time that properties were read for the RTView process specified by the selected **Connection**.

<b>Property Files</b> (table)	List of all properties files that were read by the RTView process specified by the selected <b>Connection</b> in the order they were read.
<b>Property Filters</b> (table)	List of all filters that are applied to the properties.
<b>Property Groups</b>	List of all property groups that are applied to the properties. <b>Property Groups</b> are only used when reading properties from a database.

## System Properties

This display shows the System properties for the RTView process specified by the selected Connection.

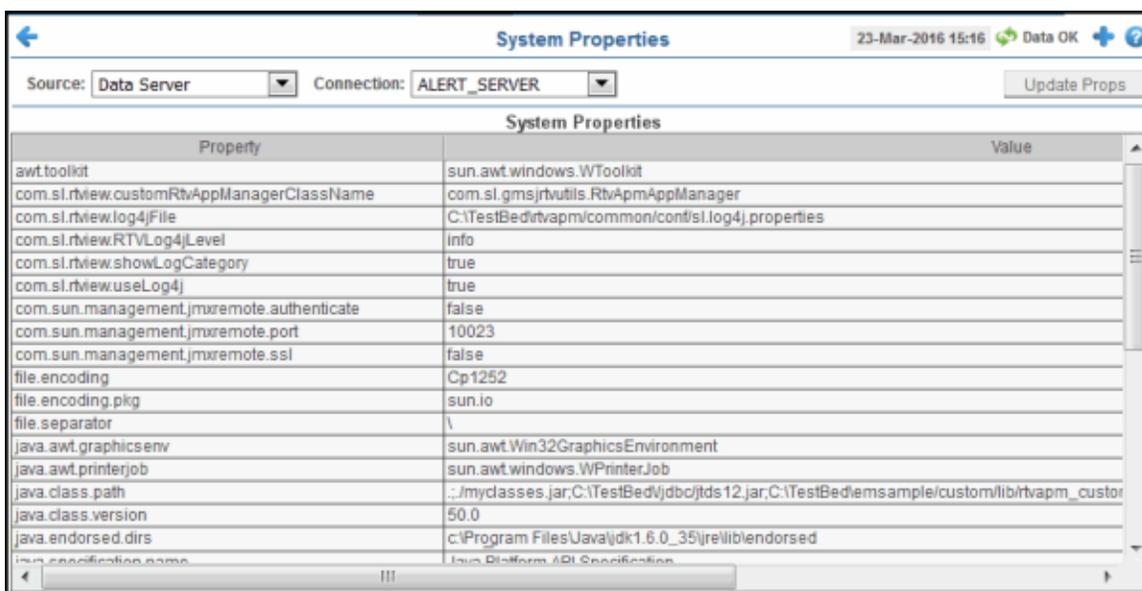
Select the **Source** of the connection to the RTView process for which you want to see property information. Options are:

- **Data Server:** If the RTView process is a Data Server and the Thin Client has a defined Data Server connection for it, choose this option and select the name of the Data Server in the Connection field.
- **Local JMX Connection:** Select this option if the Thin Client has a defined JMX Connection to the RTView process.
- **RTVMGR JMX Connection:** Select this option if the RTView has a defined JMX Connection to the RTView process.

Select the **Connection** to the RTView process for which you want to see property information. Options available depend on your setup. For example, **RTVMGR** is only visible when the **Source** is **RTVMGR JMX Connection** and you have multiple RTView s. You can then select an RTView that has a defined JMX Connection to the RTView process for which you want to see property information.

Click **Update Props** to have the RTView process specified by the selected Connection re-read all properties files and database properties. Note that most non-connection properties do NOT support updates. See the "[Properties Descriptions](#)" display to find out if a specific property supports updates.

Right-click/**Export** to create a PDF image of the display. Click Sort  to order column data.



Property	Value
awt.toolkit	sun.awt.windows.WToolkit
com.sl.rtvew.customRtvAppManagerClassName	com.sl.gmsjrtvutils.RtvApmAppManager
com.sl.rtvew.log4jFile	C:\TestBed\rtvapm\common\conf\sl.log4j.properties
com.sl.rtvew.RTVLog4jLevel	Info
com.sl.rtvew.showLogCategory	true
com.sl.rtvew.useLog4j	true
com.sun.management.jmxremote.authenticate	false
com.sun.management.jmxremote.port	10023
com.sun.management.jmxremote.ssl	false
file.encoding	Cp1252
file.encoding.pkg	sun.io
file.separator	\
java.awt.graphicsenv	sun.awt.Win32GraphicsEnvironment
java.awt.printerjob	sun.awt.windows.WPrinterJob
java.class.path	..\myclasses.jar;C:\TestBed\jdbco\jds12.jar;C:\TestBed\itemsample\custom\lib\rtvapm_custom
java.class.version	50.0
java.endorsed.dirs	c:\Program Files\Java\jdk1.6.0_35\jre\lib\endorsed
java.specification.name	Java Platform API Specification

#### Title Bar (possible features are):

-   Open the previous and upper display.
-  Open an instance of this display in a new window.
-  Open the online help page for this display.
- open commonly accessed displays.

-  Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

#### Filter By:

- Source:** Select the **Source** of the connection to the RTView process for which you want to see property information.
- Connection:** Select the **Connection** to the RTView process for which you want to see property information.

#### Fields and Data

This display includes:

- Update Props** Click to have the RTView process specified by the selected Connection re-read all properties files and database properties. Note that most non-connection properties do NOT support updates. Use the "Properties Descriptions" display to see if a specific property supports updates.
- System Properties (table)** List of all system properties for the RTView process specified by the selected **Connection**.
- Property** The name of the property.
- Value** The property setting.

### Applied Properties

This display shows all properties that were applied to the RTView process specified by the selected **Connection**.

There are several reasons a property specified in a properties file might not be applied to an RTView process:

- the filter doesn't match.
- it was overridden in another property file.
- it was specified in a file that is not used by the RTView process.
- it was a property that is not supported in that RTView process (ex, a builder specific property would not be applied to a data server process).

You can filter the **Applied Properties** table using the **Filter Column** and **Field Value** fields. The **Clear Filter** button clears the filter. Double-click on a row in the table to drill down to the ["All Properties"](#) display filtered by the **Property Name** for that row.

---

**Note:** The double-click feature is not supported on iPad. iPad users can access the ["All Properties"](#) display from the navigation tree.

---

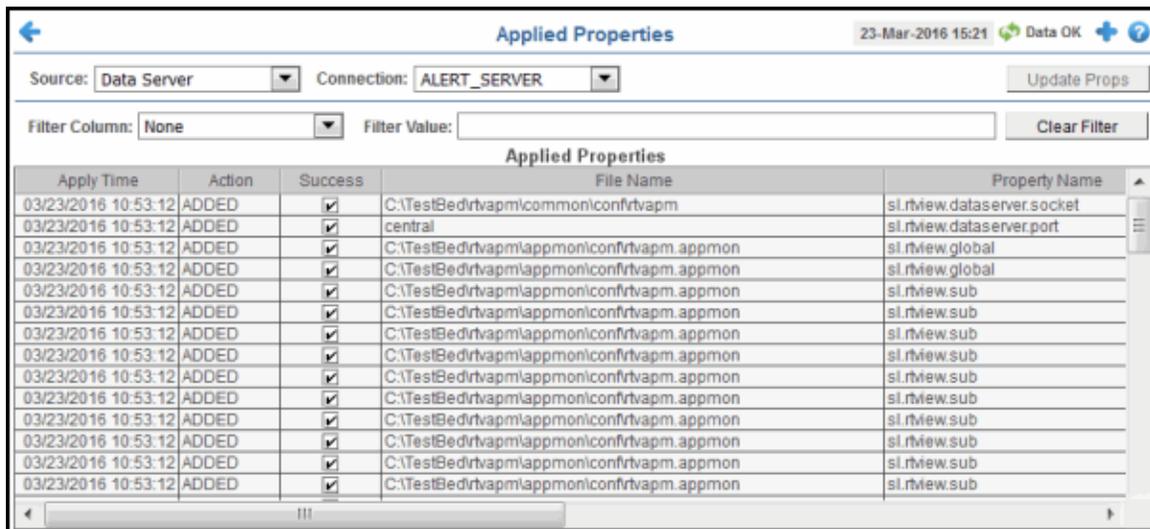
Select the **Source** of the connection to the RTView process for which you want to see property information. Options are:

- **Data Server:** If the RTView process is a Data Server and the Thin Client has a defined Data Server connection for it, choose this option and select the name of the Data Server in the Connection field.
- **Local JMX Connection:** Select this option if the Thin Client has a defined JMX Connection to the RTView process.
- **RTVMGR JMX Connection:** Select this option if the RTView has a defined JMX Connection to the RTView process.

Select the **Connection** to the RTView process for which you want to see property information. Options available depend on your setup. For example, **RTVMGR** is only visible when the **Source** is **RTVMGR JMX Connection** and you have multiple RTView s. You can then select an RTView that has a defined JMX Connection to the RTView process for which you want to see property information.

Click **Update Props** to have the RTView process specified by the selected Connection re-read all properties files and database properties. Note that most non-connection properties do NOT support updates. See the ["Properties Descriptions"](#) display to find out if a specific property supports updates.

Right-click/**Export** to create a PDF image of the display. Click Sort  to order column data.



Apply Time	Action	Success	File Name	Property Name
03/23/2016 10:53:12	ADDED	<input checked="" type="checkbox"/>	C:\TestBed\rtvapl\common\conf\rtvapl\...	sl.rtvew.dataserver.socket
03/23/2016 10:53:12	ADDED	<input checked="" type="checkbox"/>	central	sl.rtvew.dataserver.port
03/23/2016 10:53:12	ADDED	<input checked="" type="checkbox"/>	C:\TestBed\rtvapl\appmon\conf\rtvapl\appmon	sl.rtvew.global
03/23/2016 10:53:12	ADDED	<input checked="" type="checkbox"/>	C:\TestBed\rtvapl\appmon\conf\rtvapl\appmon	sl.rtvew.global
03/23/2016 10:53:12	ADDED	<input checked="" type="checkbox"/>	C:\TestBed\rtvapl\appmon\conf\rtvapl\appmon	sl.rtvew.sub
03/23/2016 10:53:12	ADDED	<input checked="" type="checkbox"/>	C:\TestBed\rtvapl\appmon\conf\rtvapl\appmon	sl.rtvew.sub
03/23/2016 10:53:12	ADDED	<input checked="" type="checkbox"/>	C:\TestBed\rtvapl\appmon\conf\rtvapl\appmon	sl.rtvew.sub
03/23/2016 10:53:12	ADDED	<input checked="" type="checkbox"/>	C:\TestBed\rtvapl\appmon\conf\rtvapl\appmon	sl.rtvew.sub
03/23/2016 10:53:12	ADDED	<input checked="" type="checkbox"/>	C:\TestBed\rtvapl\appmon\conf\rtvapl\appmon	sl.rtvew.sub
03/23/2016 10:53:12	ADDED	<input checked="" type="checkbox"/>	C:\TestBed\rtvapl\appmon\conf\rtvapl\appmon	sl.rtvew.sub
03/23/2016 10:53:12	ADDED	<input checked="" type="checkbox"/>	C:\TestBed\rtvapl\appmon\conf\rtvapl\appmon	sl.rtvew.sub
03/23/2016 10:53:12	ADDED	<input checked="" type="checkbox"/>	C:\TestBed\rtvapl\appmon\conf\rtvapl\appmon	sl.rtvew.sub
03/23/2016 10:53:12	ADDED	<input checked="" type="checkbox"/>	C:\TestBed\rtvapl\appmon\conf\rtvapl\appmon	sl.rtvew.sub

#### Title Bar (possible features are):

-   Open the previous and upper display.
-  Open an instance of this display in a new window.
-  Open the online help page for this display.
-   open commonly accessed displays.

-  **Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
-  **23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

#### Filter By:

- Source:** Select the **Source** of the connection to the RTView process for which you want to see property information.
- Connection:** Select the **Connection** to the RTView process for which you want to see property information.

#### Fields and Data

This display includes:

- Update Props** Click to have the RTView process specified by the selected Connection re-read all properties files and database properties. Note that most non-connection properties do NOT support updates. Use the "Properties Descriptions" display to see if a specific property supports updates.
- Filter Column:** Select a column to filter the **Applied Properties** table.
- Filter Value:** Enter a string to filter the **Applied Properties** table.
- Clear Filter** Clears the filter.
- Apply Time** The last time this property was applied.

#### Applied Properties (table)

<b>Action</b>	Describes what occurred at <b>Apply Time</b> . <ul style="list-style-type: none"> <li>• <b>ADDED</b>: Property was added.</li> <li>• <b>REMOVED</b>: Property was removed.</li> <li>• <b>CHANGED</b>: Property was modified.</li> </ul>
<b>Success</b>	When the box is checked the <b>Action</b> was successful.
<b>File Name</b>	The source of this property. For properties read from a database this value is <b>database</b> .
<b>Property Name</b>	The name of the property after the property filter has been applied.
<b>Property Value</b>	The value of the property.
<b>Handler</b>	The RTView Handler that uses this property.

## All Properties

This display shows all properties that were read from the properties files and database regardless of whether or not the RTView process uses them. There are several reasons a property specified in a properties file might not be applied to an RTView process:

- the filter doesn't match.
- it was overridden in another property file.
- it was specified in a file that is not used by the RTView process.
- it was a property that is not supported in that RTView process (ex, a builder specific property would not be applied to a data server process).

You can filter the **All Properties** table using the **Filter Column** and **Field Value** fields. The **Clear Filter** button clears the filter. Double-click on a row in the table to drill down to the "[All Properties](#)" display filtered by the **Property Name** for that row.

Select the **Source** of the connection to the RTView process for which you want to see property information. Options are:

- **Data Server**: If the RTView process is a Data Server and the Thin Client has a defined Data Server connection for it, choose this option and select the name of the Data Server in the Connection field.
- **Local JMX Connection**: Select this option if the Thin Client has a defined JMX Connection to the RTView process.
- **RTVMGR JMX Connection**: Select this option if the RTView has a defined JMX Connection to the RTView process.

Select the **Connection** to the RTView process for which you want to see property information. Options available depend on your setup. For example, **RTVMGR** is only visible when the **Source** is **RTVMGR JMX Connection** and you have multiple RTView s. You can then select an RTView that has a defined JMX Connection to the RTView process for which you want to see property information.

Click **Update Props** to have the RTView process specified by the selected **Connection** re-read all properties files and database properties. Note that most non-connection properties do NOT support updates. See the “[Properties Descriptions](#)” display to find out if a specific property supports updates.

Order	File Name	Property Name	Property Value
0	C:\TestBed\rtv\apm\common\conf\rtv\apm	sl.rtvview.cp	C:\TestBed\rtv\apm\rtview\lib\rtvssa.jar
1	C:\TestBed\rtv\apm\common\conf\rtv\apm	sl.rtvview.cp	C:\TestBed\rtv\apm\common\lib\rtv\apm_common.jar
2	C:\TestBed\rtv\apm\common\conf\rtv\apm	sl.rtvview.cp	C:\TestBed\rtv\apm\common\lib\gms\rtvutils.jar
3	C:\TestBed\rtv\apm\common\conf\rtv\apm	sl.rtvview.cp	C:\TestBed\rtv\apm\rtview\lib\rtvdebug.jar
4	C:\TestBed\rtv\apm\common\conf\rtv\apm	sl.rtvview.jvm	-Xmx256m
5	C:\TestBed\rtv\apm\common\conf\rtv\apm	sl.rtvview.jvm	-Xms128m
6	C:\TestBed\rtv\apm\common\conf\rtv\apm	sl.rtvview.cmd_line	-nolibco
7	C:\TestBed\rtv\apm\common\conf\rtv\apm	sl.rtvview.stylesheet	rtv_darkstyles.rtv_flat.rtv_html5
8	C:\TestBed\rtv\apm\common\conf\rtv\apm	sl.rtvview.sql.dbretry	40000
9	C:\TestBed\rtv\apm\common\conf\rtv\apm	sl.rtvview.global	rtv_global_vars.rtv
10	C:\TestBed\rtv\apm\common\conf\rtv\apm	sl.rtvview.global	rtv_global_trendrange.rtv
11	C:\TestBed\rtv\apm\common\conf\rtv\apm	sl.rtvview.xml.xmlsource	rtv_constants.xml 0 rtv_constants.xml 0 1
12	C:\TestBed\rtv\apm\common\conf\rtv\apm	sl.rtvview.jmx.jmxconn	local -- 'URL_local' -- false
13	C:\TestBed\rtv\apm\common\conf\rtv\apm	sl.rtvview.dsenable	jmx
14	C:\TestBed\rtv\apm\common\conf\rtv\apm	sl.rtvview.jmx.jmx_metrics_period	10000
15	C:\TestBed\rtv\apm\common\conf\rtv\apm	sl.rtvview.jmx.jmx_minreconnecttime	30
16	C:\TestBed\rtv\apm\common\conf\rtv\apm	sl.rtvview.jmx.jmx_mbeans_change_dyn	false
17	C:\TestBed\rtv\apm\common\conf\rtv\apm	sl.rtvview.jmx.jmxdsShowConnectionOnl	true
18	C:\TestBed\rtv\apm\common\conf\rtv\apm	sl.rtvview.jvm	-Dcom.sl.rtvview.customRtvAppManagerClassName=com.

#### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

- Order** The order in which this property was read. For properties that support a single value that are specified multiple times, the one with the highest Order value will be applied.
- File Name** The source of this property. It will be database for properties read from a database.
- Property Name** The name of the property after the property filter has been applied.
- Property Value** The value of the property.
- Original Property Name** The name of the property before the property filter was applied. This will match the literal property string in your properties file.

#### Filter By:

- Source:** Select the **Source** of the connection to the RTView process for which you want to see property information.

**Connection:** Select the **Connection** to the RTView process for which you want to see property information.

### Fields and Data

This display includes:

<b>Update Props</b>	Click to have the RTView process specified by the selected Connection re-read all properties files and database properties. Note that most non-connection properties do NOT support updates. Use the <a href="#">"Properties Descriptions"</a> display to see if a specific property supports updates.	
<b>Filter Column:</b>	Select a column to filter the <b>Applied Properties</b> table.	
<b>Filter Value:</b>	Enter a string to filter the <b>Applied Properties</b> table.	
<b>Clear Filter</b>	Clears the filter.	
<b>All Properties (table)</b>	<b>Order</b>	The order in which this property was read. For properties that support a single value that are specified multiple times, the one with the highest Order value will be applied.
	<b>File Name</b>	The source of this property. It will be database for properties read from a database.
	<b>Property Name</b>	The name of the property after the property filter has been applied.
	<b>Property Value</b>	The value of the property.
	<b>Original Property Name</b>	The name of the property before the property filter was applied. This will match the literal property string in your properties file.

## Properties Descriptions

This display shows one row for each property that is supported for the RTView process specified by the selected Connection.

Select the **Source** of the connection to the RTView process for which you want to see property information. Options are:

- **Data Server:** If the RTView process is a Data Server and the Thin Client has a defined Data Server connection for it, choose this option and select the name of the Data Server in the Connection field.
- **Local JMX Connection:** Select this option if the Thin Client has a defined JMX Connection to the RTView process.
- **RTVMGR JMX Connection:** Select this option if the RTView has a defined JMX Connection to the RTView process.

Select the **Connection** to the RTView process for which you want to see property information. Options available depend on your setup. For example, **RTVMGR** is only visible when the **Source** is **RTVMGR JMX Connection** and you have multiple RTView s. You can then select an RTView that has a defined JMX Connection to the RTView process for which you want to see property information.

Property	Multi	Updates	Handler	Deprecated
si.rview.alert.actionauditdataserver	<input type="checkbox"/>	<input type="checkbox"/>	Alert Data Source	<input type="checkbox"/>
si.rview.alert.actionauditdb	<input type="checkbox"/>	<input type="checkbox"/>	Alert Data Source	<input type="checkbox"/>
si.rview.alert.actionaudittable	<input type="checkbox"/>	<input type="checkbox"/>	Alert Data Source	<input type="checkbox"/>
si.rview.alert.alertclearedcommand	<input type="checkbox"/>	<input type="checkbox"/>	Alert Data Source	<input type="checkbox"/>
si.rview.alert.alertcleartime	<input type="checkbox"/>	<input type="checkbox"/>	Alert Data Source	<input type="checkbox"/>
si.rview.alert.alertcommand	<input type="checkbox"/>	<input type="checkbox"/>	Alert Data Source	<input type="checkbox"/>
si.rview.alert.alertinitdelay	<input type="checkbox"/>	<input type="checkbox"/>	Alert Data Source	<input type="checkbox"/>
si.rview.alert.cleansettingstable	<input type="checkbox"/>	<input type="checkbox"/>	Alert Data Source	<input type="checkbox"/>
si.rview.alert.commentcommand	<input type="checkbox"/>	<input type="checkbox"/>	Alert Data Source	<input type="checkbox"/>
si.rview.alert.commentlimit	<input type="checkbox"/>	<input type="checkbox"/>	Alert Data Source	<input type="checkbox"/>
si.rview.alert.config	<input checked="" type="checkbox"/>	<input type="checkbox"/>	Alert Data Source	<input type="checkbox"/>
si.rview.alert.createDbTables	<input type="checkbox"/>	<input type="checkbox"/>	Alert Data Source	<input type="checkbox"/>
si.rview.alert.customAlertActionHandlerClass	<input type="checkbox"/>	<input type="checkbox"/>	Alert Data Source	<input type="checkbox"/>
si.rview.alert.custom_alertdef_prop	<input checked="" type="checkbox"/>	<input type="checkbox"/>	Alert Data Source	<input type="checkbox"/>
si.rview.alert.custom_event_attr	<input checked="" type="checkbox"/>	<input type="checkbox"/>	Alert Data Source	<input type="checkbox"/>
si.rview.alert.enableactionaudit	<input type="checkbox"/>	<input type="checkbox"/>	Alert Data Source	<input type="checkbox"/>
si.rview.alert.enablebuffer	<input type="checkbox"/>	<input type="checkbox"/>	Alert Data Source	<input type="checkbox"/>
si.rview.alert.enabled	<input type="checkbox"/>	<input type="checkbox"/>	Alert Data Source	<input type="checkbox"/>
si.rview.alert.enablessa	<input type="checkbox"/>	<input type="checkbox"/>	Alert Data Source	<input type="checkbox"/>
si.rview.alert.exitOnPersistInitFailed	<input type="checkbox"/>	<input type="checkbox"/>	Alert Data Source	<input type="checkbox"/>
si.rview.alert.history	<input type="checkbox"/>	<input type="checkbox"/>	Alert Data Source	<input type="checkbox"/>
si.rview.alert.lutupdatesnewdata	<input type="checkbox"/>	<input type="checkbox"/>	Alert Data Source	<input type="checkbox"/>

#### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

#### Filter By:

- Source:** Select the **Source** of the connection to the RTView process for which you want to see property information.
- Connection:** Select the **Connection** to the RTView process for which you want to see property information.

#### Fields and Data

This display includes:

- Update Props** Click to have the RTView process specified by the selected Connection re-read all properties files and database properties. Note that most non-connection properties do NOT support updates. Use the ["Properties Descriptions"](#) display to see if a specific property supports updates.

#### All Properties (table)

- Property** The name of the property

<b>Multi</b>	Box is checked if this property supports multiple values.
<b>Updates</b>	Box is checked if this property supports updates.
<b>Handler</b>	The name of the RTView Handler that uses this property.
<b>Deprecated</b>	Box is checked if this property is deprecated.
<b>Deprecation Info</b>	If the property is deprecated, this lists the currently supported property to use instead.

## Diagram Views

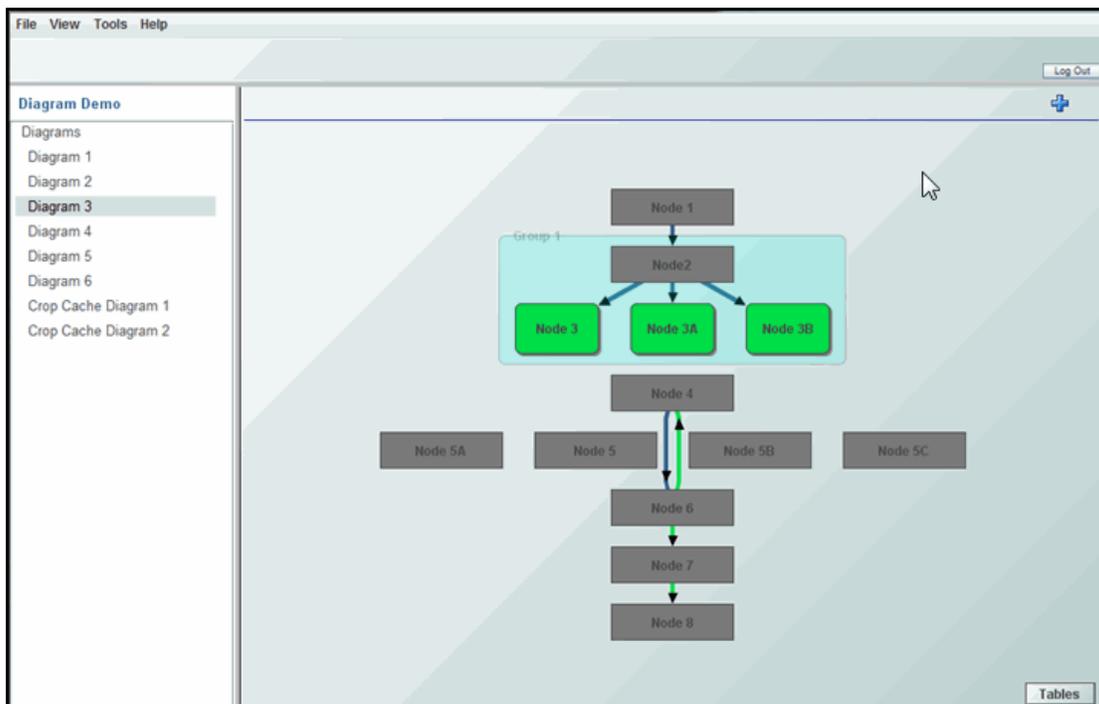
The RTView Enterprise comes with the Diagram Generator, a feature that auto-creates diagram displays which mirror your system components and hierarchy.

When you monitor applications with complex architectures, it is often very beneficial to visualize the health of individual components within the context of the application hierarchy. This allows you to understand the complete architecture which is supporting the application as well as understand how individual components may affect the behavior of other components. RTView Enterprise allows you to manually construct such views, but at times this might become too time consuming to maintain if there are many applications to model or the architecture is constantly changing. The Diagram Generator is a feature which allows for the automatic generation of these application diagrams using application meta-data, without having to manually construct them.

The Diagram Generator, located under **CUSTOM Tab/Diagram Views**, is comprised of several displays which you use to create your diagram displays.

This section includes the following Diagram Generator instructions and displays:

- "Steps to Create a Diagram Display"
- "Create an Object Template Display"
- "Node Administration Display"
- "Link Admin Display"
- "Diagram Properties Admin Display"
- "Add Diagrams to your Project"
- "View Diagram Displays"
- "Optional Diagram Display Customizations"



The Diagram Generator feature uses a database table of your nodes and database table of your links to create the diagrams. For example, the diagram below was generated from the tables shown next to it. The order of the nodes in the table controls the order of the nodes in the diagram.

The screenshot shows the RTView Enterprise Monitor interface. The main window displays a 'Sample Diagram' with the following structure:

```

graph TD
    Node1[Node 1] --> Node2[Node 2]
    Node2 --> Node3B[Node 3B]
    Node3A[Node 3A] <--> Node3B
    Node3B <--> Node3C[Node 3C]
    Node3B --> Node4[Node 4]
    subgraph Group1 [Group 1]
        Node4
        Node5[Node 5]
    end
  
```

Below the diagram are three data tables:

Nodes			
ID	Type	Properties	Layout
G1	Group DefaultGroup	label=Group 1	
N1	DefaultNode	label=Node 1~dd=rtv_diagram_test2.rtv~mouseOverText=dd to diagram	
N2	DefaultNode	label=Node2	group=G1
N3	GreenNode	label=Node 3	tier=SL_1~group=G1
N3A	GreenNode	label=Node 3A	tier=SL_1~group=G1
N3B	GreenNode	label=Node 3B	tier=SL_1~group=G1
N4	DefaultNode	label=Node 4	
N5A	DefaultNode	label=Node 5A	tier=SL_2

Links				
ID	Type	Properties	Start Node	End Node
L1	DefaultLink		N1	N2
L2	DefaultLink		N2	N3
L3	DefaultLink		N2	N3A
L4	DefaultLink		N2	N3B
L5	DefaultLink		N4	N6
L6	GreenLink		N6	N4
L7	GreenLink		N6	N7
L8	GreenLink		N7	N8

Diagram Props	
Name	Value
hGap	30
objectTemplateDisplay	test_clone_source.rtv
orientation	VERTICAL
hAlignment	CENTER
vAlignment	CENTER
margins	40,10,10,10
wrapDiagram	false
wrapSpacing	20

### Diagram Generator Demo

You can view a demo of the Diagram Generator, located under **CUSTOM/Sample Diagram**. This demo shows a simple diagram as well as the UI used to construct diagrams. The icons in the sample diagram come from a sample object template display. When creating diagrams for your project, you will create an object template display with icons that are appropriate for the process you wish to diagram.

### Steps to Create a Diagram Display

To create a custom diagram display using the Diagram Generator:

1. If you are using EM 3.5 or earlier, see **Upgrading the** . If you are using EM 3.6+, the **emsample** project is already configured to include the Diagram Generator and no setup is required.
2. **“Create an Object Template Display”** using the Display Builder. This step is optional. The object template display serves as your palette of objects (icons, links, shapes, and so forth) for building your diagrams. A sample object template display is built-in that can be used for simple diagrams and to demo the Diagram Generator. When creating diagrams for your project, you should create an object template display with icons that are appropriate for the process you wish to diagram. If you are going to create a custom object template display, you must create it and add it to the Diagram Properties before defining nodes and links for your diagrams.
3. In the RTView Enterprise, open the **“Node Administration Display”** display, located under **CUSTOM Tab/Diagram Views/Diagram Admin** and add nodes to your diagram.
4. Open the **“Link Admin Display”** display and add links to your diagram.
5. Open the **“Diagram Properties Admin Display”** display and format the layout of your diagram display.
6. **“Add Diagrams to your Project”** to publish the diagram display.
7. **“View Diagram Displays”** to confirm settings.
8. **“Optional Diagram Display Customizations”**: These customizations are not required.

## Create an Object Template Display

This section describes how to create an object template display using the Display Builder. The object template display serves as your palette of objects (icons, links, shapes, and so forth) that you use to build your diagrams.

### Assumptions:

- You have familiarized yourself with the **“Diagram Generator Demo”**.
  - You are familiar with using the Display Builder.
1. Create a display in the Display Builder that contains all the icons you want to use for the nodes, links and icons in your diagram displays.
  2. For each object, specify a user friendly name in the **objName** field (it must be alpha-numeric but can contain under-bar). The **objName** field will be referenced in the **Type** field when you add nodes and links to your diagram. You will be able to override any of the properties by specifying them in the **Properties** field.
  3. Save this file to the **RTViewEnterprise/emsample/servers/central** directory.
  4. Use the **“Diagram Properties Admin Display”** to set the **objectTemplateDisplay** property to the name of this file.

For assistance, contact Technical Support.

Proceed to **“Node Administration Display”**.

## Node Administration Display

Use this Diagram Generator display to create new diagrams and add or edit nodes in existing diagrams. Diagram definitions are stored in the DIAGRAMS database.

To define a new diagram, enter a new diagram name in the **Diagram** field, then fill in the fields in the **Enter Values** section to define the first node in the diagram.

To add a new node to a diagram, enter the name of the diagram to which you want to add the node, then fill in the fields in the **Enter Values** section, described below, and click **Add Node**.

To edit an existing node, select it in the table to populate the **Enter Values** fields at the bottom of the display, make your changes and click **Update Node**.

The screenshot shows the 'Node Table Administration' interface in RTView Enterprise Monitor. The top navigation bar includes 'SERVICE TREE', 'SERVICE VIEWS', 'COMPONENTS', 'ALERTS', 'ADMIN', and 'CUSTOM'. The main content area features a 'Nodes' dropdown menu and a 'Diagram Filter' set to 'All'. A table lists nodes with columns for Order, Diagram, ID, Type, Layout, and Properties. Below the table is the 'Enter Values' section with input fields for Diagram, ID, Layout, Type, and Properties, along with buttons for 'Update Node', 'Insert Node', 'Delete Node', 'Cancel', 'Move Node Up', 'Move Node Down', and 'Preview'.

Order	Diagram	ID	Type	Layout	Properties
0	sample_diagram	N1	DefaultNode		label=Node 1
1	sample_diagram	N2	DefaultNode		label=Node 2
2	sample_diagram	N3	GreenNode	tier=3	label=Node 3A
3	sample_diagram	N4	GreenNode	tier=3	label=Node 3B
4	sample_diagram	N5	GreenNode	tier=3	label=Node 3C
5	sample_diagram	N6	DefaultNode	group=G1	label=Node 4
6	sample_diagram	N7	DefaultNode	group=G1	label=Node 5
7	sample_diagram	G1	Group DefaultGroup		label=Group 1

For each node in your diagram display, fill in the following fields:

- Diagram** Required. The name of the diagram. **Note:** The diagram name GLOBAL is reserved for global properties in Diagram Props. The value in this column will be used to identify this diagram when you add this diagram to the navigation tree.
- ID** Required. Must be unique within the diagram across nodes and links. Use this value in the Link Node1 and Node2 fields to refer to this node. The value must be alpha-numeric but may contain underbars (\_).
- Layout** Optional. One or more layout options for the icon. All icons with the same tier=value will be positioned in a single tier (row if the diagram property orientation=VERTICAL, column if orientation=HORIZONTAL). Each tier is centered along the diagram's centerline unless only one node in a tier has a link to another tier. In that case the node with the link is centered on the diagram's centerline.
- All icons with the same group=value will have a group object drawn behind them. The group value must be the ID of a node whose type is defined as a Group.
- To specify both a tier and a group, separate them with a ~. For example:  
**tier=T1~group=G1**

<b>Type</b>	Required. The name of the object in the objectTemplateDisplay file to use as the icon for this node. When you select a type from the list, you will see a preview of it to the right of the Type field.
<b>Is Group</b>	Optional. Check to specify that this node is a Group. Groups are only drawn if at least one node references them in their Layout field. They are drawn behind the nodes that reference them and their extent is set to the combined extent of all nodes that reference them. In wrapped diagrams, if the nodes in the group break across multiple tiers, the group object will be broken across the tiers as well.
<b>Properties</b>	<p>Optional. One or more properties to set on the node icon delimited by ~. Syntax is <b>propName=propVal~propName2=otherPropVal</b>.</p> <p>Note that property values must be specified as they are saved in <b>.rtv</b> files, which is not necessarily the same as they are shown in the <b>Object Properties</b> dialog in the Display Builder.</p> <p>In addition to properties on the RTView object, you can also specify hGap or vGap to override the diagram property hGap or vGap for this object. The hGap is applied to the left of an object and the vGap is applied above an object.</p> <p>For example, you must specify the font index instead of the font name for font properties, and the color index instead of the color for color properties.</p>

Use the following buttons to save changes to the database and to preview the diagram display (after the changes have been saved to the database):

<b>Update Node</b>	Save changes to the selected node to the database. This is only enabled if the selected node is already in the database.
<b>Insert Node</b>	Insert a new node to the database. This is only enabled if the selected node is not in the database.
<b>Delete Node</b>	Delete the selected node from the database. This is only enabled if the selected node is already in the database.
<b>Cancel</b>	Clear the Enter Values fields.
<b>Move Node Up</b>	Move the selected node up in the diagram. Nodes are laid out in the diagram according to their order.
<b>Move Node Down</b>	Move the selected node down in the diagram. Nodes are laid out in the diagram according to their order.
<b>Preview</b>	Open a window showing the selected diagram as it is saved in the database. Changes to the diagram will not update an open preview window. To update the diagram in the preview window, close and reopen the window.

## Link Admin Display

Use this display to add or edit links in your diagrams. To add a link, enter the name of the diagram containing the nodes you want to link, then fill in the fields below for each link you want to add. To edit an existing link, select it in the table.

The screenshot shows the 'Link Table Administration' window in RTView Enterprise Monitor. The window has a navigation pane on the left with 'Link Admin' selected. The main area contains a table of links and a form to add or edit links.

Diagram	ID	Type	Node 1	Node 2	
sample_diag	L1	DefaultLink	N1	N2	
sample_diag	L2	DefaultLink	N2	N4	
sample_diag	L3	DefaultLink	N4	N6	
sample_diag	L4	GreenLink	N3	N4	arrow1VisFlag=1
sample_diag	L5	GreenLink	N4	N5	arrow1VisFlag=1

Below the table is a form to enter values for a new link:

Enter Values

Diagram:  ID:  Type:

Node 1:  Node 2:

Properties:

Buttons: Update Link, Insert Link, Delete Link, Cancel, Preview

- Diagram** Required. The name of the diagram. **Note:** The diagram name GLOBAL is reserved for global properties in Diagram Props.
- ID** Required. Must be unique within the diagram across the node and link tables. Value must be alpha-numeric but may contain under-bars (\_).
- Type** Required. The name of the link object in the objectTemplateDisplay to use for this link.
- Node 1** Required. The ID of the start node for the link.
- Node 2** Required. The ID of the end node for the link.
- Properties** Optional. One or more properties to set on the icon delimited by ~. Syntax is propName=propVal~propName2=otherPropVal

Use the following buttons to save link changes to the database and to preview the diagram display (after the changes have been applied to the database):

- Update Link** Save changes to the selected link to the database. This is only enabled if the selected link is already in the database.
- Insert Link** Insert a new link to the database. This is only enabled if the selected link is not in the database.
- Delete Link** Delete the selected link from the database. This is only enabled if the selected link is already in the database.

**Cancel** Clear the **Enter Values** fields.

**Preview** Opens a window showing the selected diagram as it is saved in the database. Changes to the diagram will not update the preview window. To update the preview, close and reopen it.

## Diagram Properties Admin Display

Use this Diagram Generator display to configure “Diagram Properties” for your diagrams. Diagram properties are settings that are applied to the diagrams as a whole, such as orientation, alignment and spacing. Properties that use GLOBAL for the diagram name are applied to all diagrams. You can override a diagram property for a single diagram by using the name of that diagram in the **Diagram** field.

The screenshot shows the 'Diagram Properties Administration' window in RTView Enterprise Monitor. The top navigation bar includes 'SERVICE TREE', 'SERVICE VIEWS', 'COMPONENTS', 'ALERTS', 'ADMIN', and 'CUSTOM'. The main content area features a 'Diagram Filter' dropdown set to 'All' and a table with the following data:

Diagram	Property Name	Property Value
GLOBAL	hAlignment	CENTER
GLOBAL	orientation	VERTICAL
GLOBAL	vAlignment	CENTER

Below the table is the 'Enter Values' form with the following fields:

- Diagram:** GLOBAL
- Property:** vAlignment
- Value:** CENTER

Buttons at the bottom of the form include 'Update Property', 'Insert Property', 'Delete Property', and 'Cancel'.

**Diagram** Required. The name of the diagram or GLOBAL if it should be applied to all diagrams.

**Property** Required. The name of the property.

**Value** Required. The value of the property.

Use the following buttons to save diagram property changes to the database:

**Update Property** Save changes to the selected property to the database. This is only enabled if the selected property is already in the database.

**Insert Property** Insert a new property to the database. This is only enabled if the selected property is not in the database.

**Delete Property** Delete the selected property from the database. This is only enabled if the selected link is already in the database.

**Cancel** Clear the Enter Values fields.

## Diagram Properties

Use the following properties in the **Property/Value** fields in the ["Diagram Properties Admin Display"](#) display to configure diagram display.

Property Name	Description
<b>hGap</b>	Horizontal space between nodes in pixels. This can be overridden per-node in the node properties. Default is 28.
<b>vGap</b>	Vertical space between nodes in pixels. This can be overridden per-node in the node properties. Default is 22.
<b>objectTemplateDisplay</b>	The name of the file containing the nodes and links to use for the diagrams. The name of each node and link in this file corresponds to the name in the Type field of the node and link tables.
<b>orientation</b>	The orientation of the diagram. Default is VERTICAL which lays the nodes out in the order specified from the top of the display to the bottom with nodes in the same tier laid out left to right. HORIZONTAL lays out the nodes in the order specified from left to right with nodes in the same tier laid out top to bottom.
<b>hAlignment</b>	Controls the horizontal alignment of the diagram in the available space. Default is CENTER which centers the diagram in the available space. Options are: LEFT – Position the diagram at the left of the available space. CENTER – Position the diagram in the center of the available space. CENTERLINE – Position the centerline of the diagram in the center of the available space. This option is only valid if orientation=VERTICAL and wrapDiagram=false. If orientation=HORIZONTAL or wrapDiagram=true, CENTER will be used instead. RIGHT – Position the diagram at the right of the available space.
<b>vAlignment</b>	Controls the vertical alignment of the diagram in the available space. Default is CENTER which centers the diagram in the available space. Options are: TOP – Position the diagram at the top of the available space. CENTER – Position the diagram in the center of the available space. CENTERLINE – Position the centerline of the diagram in the center of the available space. This option is only valid if orientation=HORIZONTAL and wrapDiagram=false. If orientation=VERTICAL or wrapDiagram=true, CENTER will be used instead. Bottom – Position the diagram at the bottom of the available space.
<b>wrapDiagram</b>	If true, the diagram wraps into columns if orientation=VERTICAL or rows if orientation=HORIZONTAL. Default is false.
<b>wrapSpacing</b>	The number of pixels between columns/rows if wrapDiagram = true. Default is 20.
<b>margins</b>	Sets the minimum amount of space between the edge of the display and the diagram in pixels. You can either specify one value that will be used for all 4 margins or a comma separated list of 4 values in this order: top, left, bottom, right. Default is 40,10,10,10 which sets the top margin to 40 pixels and the left, bottom and right margins to 10 pixels.
<b>deleteSavedDiagramNodes</b>	This option applies to diagrams that were manually edited as described in <a href="#">"Edit Diagrams in the Display Builder"</a> . If true, delete any diagram nodes that were saved to the display in the Display Builder. Default is false.

Proceed to ["Add Diagrams to your Project"](#) to publish your displays.

## Add Diagrams to your Project

This section describes how to add diagrams to your EM project. After you define one or more diagrams as described in [“Node Administration Display”](#), you can add a display for each diagram to the EM navigation tree. The navigation tree entry for each diagram should look like this:

```
<node label="Diagram 1" display="rtv_diagram_cache" subs="$diagramName:diagram1 $diagramTitle:'Diagram 1'"/>
```

The value for **label** is the label to use in the navigation tree. The value for **display** is **rtv\_diagram\_cache** unless you have a custom diagram background as described in [“Customize the Diagram Background Display”](#), in which case you should use the name of that display instead.

The **subs** values are as follows:

- **\$diagramName** – Set this to the name of your diagram. This corresponds to the value in the DIAGRAM column in the database.
- **\$diagramTitle** – Set this to the value to use for the title label in the diagram display.

---

**Note:** You must use single quotes around any substitution values that contain spaces.

---

Proceed to [“View Diagram Displays”](#).

## View Diagram Displays

After you add one or more diagram displays to your EM project as described in the [“Add Diagrams to your Project”](#), open them in the navigation tree. Note that:

- Diagram definitions are only read when the display is opened. If you edit the diagram definition for an open diagram display, you must re-open the diagram display to see the changes.
- When you resize the window, the diagram display auto-resizes to fill the available space, and also positions the diagram in the available space according to the **margin** and **alignment** Diagram Properties. If you resize the window smaller than 800x576 or the area required to display the diagram (whichever is larger) scrollbars auto-appear.
- Diagrams with the **wrapDiagram** Diagram Property set to true reposition nodes to use the available space when the window is resized.

## Optional Diagram Display Customizations

This section includes:

- [“Edit Diagrams in the Display Builder”](#):
- [“Customize the Diagram Background Display”](#):
- [“Customize the Diagram Database”](#):

## Edit Diagrams in the Display Builder

To manually edit your generated diagram, use the Display Builder to open the diagram in the **RTViewEnterprise/emsample/servers/central** directory.

---

**Note:** If you created a custom diagram background display as described in [“Customize the Diagram Background Display”](#), use the name of that file instead of **rtv\_diagram\_cache** in the instructions below. Run the Display Builder in the **RTViewEnterprise/emsample/servers/central** directory as follows (where **diagramName** is the name of the diagram you want to modify):

---

```
runb_appmon -sub:$diagramName:diagramName rtv_diagram_cache
```

Edit the diagram and save the display as **rtv\_diagram\_cache\_diagramName**, replacing the **diagramName** with the name of your diagram. Update the corresponding navigation tree entry to use the new display name.

**Important:** Do NOT to save these changes to **rtv\_diagram\_cache.rtv** or these nodes will show up in all of your diagrams.

When you view this diagram, any saved diagram nodes and links that are no longer in the diagram definition will be removed and any new nodes in the diagram definition will be added to the bottom left corner. You need to position those new nodes by hand in the Display Builder. The Properties from the database will be applied to diagram nodes that were saved in the Display Builder. The diagram will still be positioned in the Display Viewer according to the **alignment** and **margin** Diagram Properties when the window is resized. However, for diagrams where **wrapDiagram** is set to true, the diagram will not be re-wrapped to fit the available space.

## Customize the Diagram Background Display

To create a custom version of the diagram background display, open **rtv\_diagram\_cache.rtv** in the Display Builder from the **central** directory of your EM project as follows:

```
runb_appmon rtv_diagram_cache
```

Modify the display and save it under a new name in the **central** directory. The name must start with **rtv\_diagram**. When adding diagram displays to the navigation tree as described in [“Add Diagrams to your Project”](#), use the name of this file instead of **rtv\_diagram\_cache**.

When modifying the display, use the following guidelines:

- Do not change the **Resize Mode**. It must be set to **Crop**.
- When you resize this display in the viewer or thin client, objects will be positioned according to their anchor properties.
- Do not remove the **dg\_include\_cache.rtv** entry from the list of included files. This file reads the diagram data and creates the data structures required to generate the diagrams.

## Customize the Diagram Database

Diagram definitions are stored in the DIAGRAM database. By default, an HSQLDB database is used. Schemas for all supported databases are provided in **RTVAPM\_HOME\dg\dbconfig**. To change to another supported database, use the schema for your database to create the diagram tables and add this property to the central properties file for your project (**central.properties** in **emsample**) replacing the user name, password, URL and driver with the appropriate information for your database:

```
ConfigCollector.sl.rtvview.sql.sqlldb=DIAGRAMS sa - jdbc:hsqldb:hsq://localhost:9099/
rtvdiagram org.hsqldb.jdbcDriver - false true
```

---

## Other - HTML Displays

The displays described in this section either have no equivalent in the Display Server version (as with “[Common/Alerts - HTML](#)” displays) or they reside in different places (as “[Common/System - HTML](#)” displays do).

RTView Enterprise - HTML includes the following Views:

- “[Common/Alerts - HTML](#)”: These displays are used for managing alerts at the component level. Unlike most RTView Enterprise displays, they are not accessed from the navigation tree. You access these from other displays.
- “[Common/System - HTML](#)”: These displays show details about your RTView Enterprise system.

## Common/Alerts - HTML

This View includes the following displays:

- “[Alerts Table by Component - HTML](#)”: Track alerts associated with CIs shown in a display.
- “[Alert Detail for Component - HTML](#)”: Investigate an alert instance and its history.
- “[Alert Configuration for Component - HTML](#)”: Refine alert threshold settings.

### Alerts Table by Component - HTML

As an alternative to the **Alerts Table**, use the **Alerts Table by Component** to track and manage all alerts that are specifically associated with the CIs shown in a display.

You access the **Alerts Table by Component** by clicking  (the alert status icon) in the title bar of other displays. The display in which you click  is the source display.

**Package** provides the technology label associated with the alerts shown. For example, **Jvm**, **Tomcat** and **Host** are the technology labels for Java Virtual Machines, Tomcat applications and servers (respectively). These labels are also correlated with the RTView solution package names (for example, the Solution Package for Host Agent). **Category** lists all alert categories related to the source display.

Use the **ACK** and **Cleared** drop-downs to filter the table by **All**, **True** or **False**.

See the **Alert Level** column icon, where:



The alert reached its ALARM LEVEL threshold in the table row.



The alert reached its WARNING LEVEL threshold in the table row.

To investigate, click:

 to open the **Alert Detail for Component** where you can see the current and historical conditions that precipitated the alert being executed.

[Go to CI](#) to open the summary display for the CI associated with the alert where you can investigate utilization metrics for the CI leading up to the alert being executed.

You can search, filter, sort and choose columns to include by clicking a column header icon (to the right of each column label) and selecting **Filter**, **Sort Ascending**, **Sort Descending** or **Columns**. Right-click on a table cell to **Export to Excel**. Use **Ctrl +** click or **Shift +** click to select multiple alerts.

With one or more alerts selected, click [Own](#) to set the alert(s) owner field, [Acknowledge](#) to acknowledge the alert(s), [Unacknowledge](#) to clear the acknowledgement on previously acknowledged alert(s), [Add Comment](#) to add a comment to the alert(s).

You must be logged in as `rtvalertmgr` or `rtvadmin` to perform the **Own**, **Ack**, **Unack**, or **Comment** actions. Otherwise, you get an error dialog.

Alerts Table by Component 02-May-2019 11:05:09 ✔ DATA OK [🔗](#) [🔔](#)

Package: Host      Category: CPU:Network:Storage      Cleared:       ACK:

Alert Count: 16

Row	Update Time	Acknowledge	Cleared	Alert Level	Alert Name	Alert Index Values	
2018-Nov-09	23:54:0				HostCpuPercentHigh	SL-DEMO;SLHOST16(sL_qa)	High V
2018-Oct-01	06:20:10				HostCpuPercentHigh	SL-DEMO;SLHOST17(sL_amx)	High A
2019-May-02	03:28:5				HostMemoryUsedHigh	SL-DEMO-LX;192.168.200.92	High V
2018-Oct-01	06:19:36				HostVirtualMemoryUsedH	SL-DEMO;SLHOST17(sL_amx)	High A
2018-Oct-01	06:18:36				HostMemoryUsedHigh	SL-DEMO;SLHOST17(sL_amx)	High V
2018-Jan-12	11:38:56				HostCpuPercentHigh	SL-DEMO-LX;192.168.200.205	High A
2019-May-02	10:40:3				HostVirtualMemoryUsedH	SL-DEMO-LX;192.168.200.42	High A
2019-Apr-25	10:19:43				HostMemoryUsedHigh	SL-DEMO;SLHOST8	High V
2018-Jun-19	09:22:23				HostCpuPercentHigh	SL-DEMO-LX;192.168.200.202	High A
2018-Nov-09	10:33:54				HostVirtualMemoryUsedH	SL-DEMO;SLHOST16(sL_qa)	High A
2018-May-01	22:45:4				HostCpuPercentHigh	SL-DEMO-LX;192.168.200.202	High A

[Alert Detail](#)   [Go to CI](#)   [Own](#)   [Acknowledge](#)   [Unacknowledge](#)  
 [Add Comment](#)   [Clear All Comments](#)

## Alert Detail for Component - HTML

Use the **Alert Detail for Component** display to investigate current and historical activity of a specific alert instance as it applies to the associated CI, and also compare against **Metric History** trends of the associated CI. A trend graph for the CI associated with the alert instance. You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.

Access the **Alert Detail for Component** display by clicking [Details](#) in the **Alerts Table** or [Alert Detail](#) in the **Alerts Table by Component** display.

The **Alert History** table at the bottom of the display contains a row of data for each time the alert instance was updated. See the alert **ID**, **Row Update Time**, **Cleared** status and **Reason**, **Owner** and the **Alert Level** column icon, where:



The alert reached its ALARM LEVEL threshold in the table row.



The alert reached its WARNING LEVEL threshold in the table row.

You can search, filter, sort and choose columns to include by clicking a column header icon (to the right of each column label) and selecting **Filter**, **Sort Ascending**, **Sort Descending** or **Columns**. Right-click on a table cell to **Export to Excel**. Use **Ctrl +** click or **Shift +** click to select multiple alerts.

To investigate, click:

[Go to CI](#)

to see utilization conditions for the CI associated with the alert in a summary display.

[Admin](#) to open the **Alert Configuration for Component** display where you can see, modify and refine alert threshold settings for that particular alert. A trend graph traces the relevant alert metric for the CI so you can adjust thresholds in real-time.

**Alert Detail For Component** ▼
02-May-2019 14:09:52 ✔ DATA OK

Alert Name: **JvmCpuPercentHigh**      Severity: 🔔      Cleared: ✔ DATA UPDATE      ACK: no

Source: Instance-1-90      Connection: CRMProducer1

Alert Time: 02-May-2019 14:00      Alert Text: High Warning Limit exceeded, current value: 63.46399696813404 limit: 50.0

[Go to CI](#)
[Own](#)
[Acknowledge](#)
[Unacknowledge](#)

[Add Comment](#)
[Admin](#)

**Metric History** Log Scale:  15 minutes 🕒

**Alert History**

ID	Row Update Time	Alert Level	Cleared	Cleared Reason	Acknowledged	Owner
937671	2019-May-02 14:09:18	🔔	✔	DATA UPDATE		
937671	2019-May-02 14:03:26	🔔				
937646	2019-May-02 13:56:27	🔔	✔	DATA UPDATE		
937646	2019-May-02 13:55:49	🔔				
937635	2019-May-02 13:51:15	🔔	✔	DATA UPDATE		
937635	2019-May-02 13:51:14	🔔				

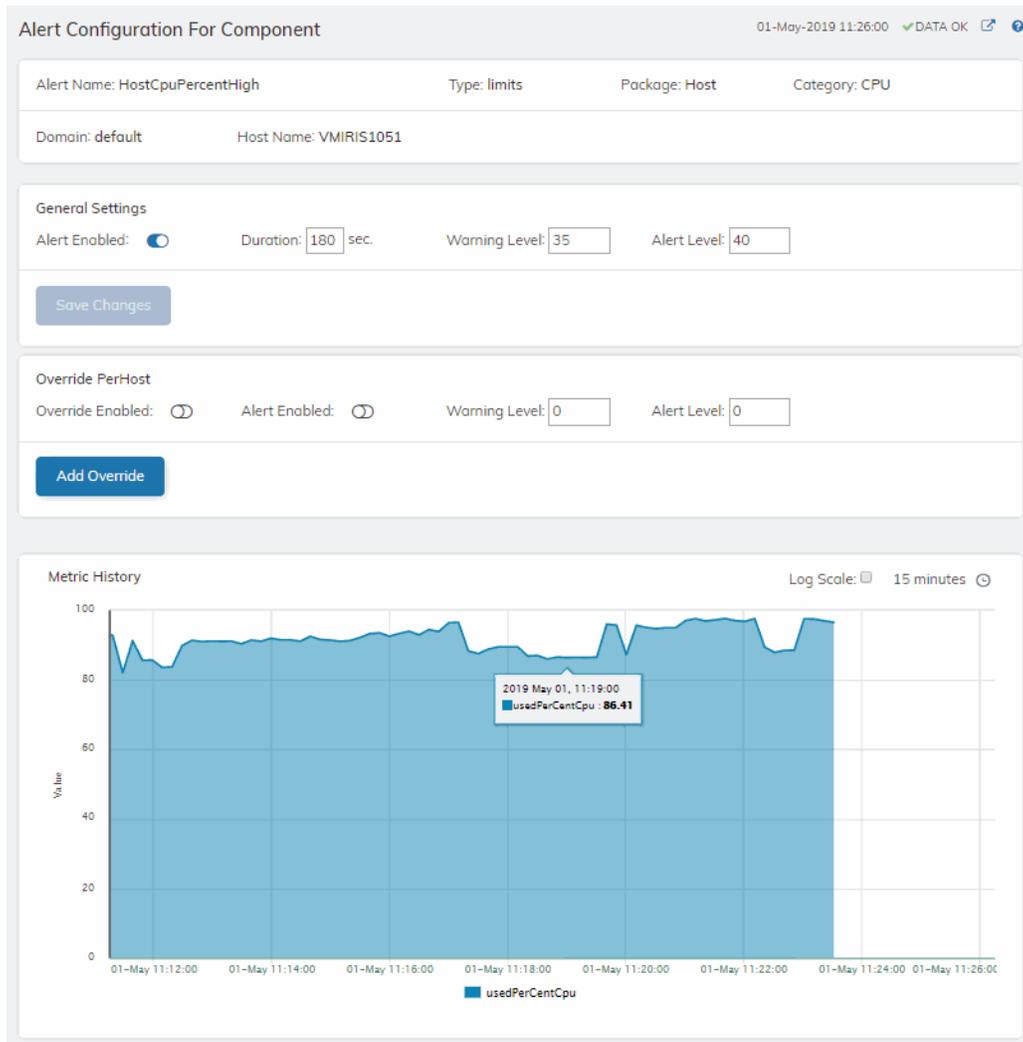
## Alert Configuration for Component - HTML

Use the **Alert Configuration for Component** display to see, modify and refine alert threshold settings for a particular alert. A trend graph traces the history of the relevant metric for this alert so you can adjust thresholds in real-time. You can also modify alert thresholds, add an override alert and toggle ON or OFF   both global and override alerts.

Access the **Alert Configuration for Component** display by clicking [Admin](#) in the **Alert Detail for Component** display.

The bottom half of the display provides a **Metric History** trend graph which traces the performance metric pertaining to the alert. You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.

You must be logged in as `rtvalertmgr` or `rtvadmin` to modify alerts.



## Common/System - HTML

RTView Enterprise - HTML includes the following displays which provide details about your RTView system setup:

- ["System Overview - HTML"](#)
- ["RTView Data Servers - HTML"](#)
- ["Data Server Summary - HTML"](#)
- ["CI Type Definitions - HTML"](#)
- ["CI Stats - HTML"](#)
- ["Cache Table - HTML"](#)

### System Overview - HTML

View the health status of RTViewCentral and all RTView DataServers connected to it. Use this display to find out if your data servers are all **Connected**. You can also get the **URL** for a DataServer, see which solution packages are installed and running on each (shown in the **Packages** field) and the number of **CI Metrics** that each data server is sending and compare the number of alerts among them.

The **Central Server** metric card shows **Configuration Server**-related metrics on the left:

- **CMDB** The number of CIs in the CMDB.
- **CIType Defs** The number of CI Type definitions.
- **Alert Defs** The number of alert settings and override definitions.

And on the right the **Central Server** metric card shows **Alert Server and Directory**-related metrics:

- **Alerts By CMDB** The number of Services in the CMDB that currently have at least one associated alert.
- **Alerts By CI** The number of CIs in the CMDB that currently have at least one associated alert.
- **RtvAlertTable** The number of currently active alerts in the system.
- **CacheMap** The number of entries currently in the directory table.

You can hover your mouse over the **Central Server** metric card to get more details.

The **Data Servers** table contains a row of data for each connected data server. You can search, filter, sort and choose columns to include by clicking a column header icon (to the right of each column label) and selecting **Filter**, **Sort Ascending**, **Sort Descending** or **Columns**. Or just click a column header to sort.

Right-click on a table cell to **Export to Excel** or **Copy Cell Value**.

**System Overview**

**Central Server** ✓

CMDB: 3442 CType Defs: 76 Alert Defs: 443	Alerts by CMDB: 76 Alerts by CI: 56 RtvAlertTable: 441 CacheMap: 2821
---	--

**Data Servers:**

Name	Connected	Alerts	CI Metrics	URL	Packages
RTV-DATA-BOOMI	✓	95	866	https://rtvdemos.sl.com/boomimon_rtvq	BOOMI-ATOM JVM
RTV-DATA-IBM	✓	0	85	https://rtvdemos.sl.com/ibmmmon_rtvque	DB2 MQ-BROKER MQ-CHANNEL MQ
RTV-DATA-INFRA	✓	0	417	https://rtvdemos.sl.com/inframmon_rtvqu	DOCKER-CONTAINER DOCKER-ENG
RTV-DATA-KAFKAMON	✓	185	142	https://rtvdemos.sl.com/kafkammon_rtvq	JVM KAFKA-BROKER KAFKA-CONSU
RTV-DATA-ORACLE	✓	0	564	https://rtvdemos.sl.com/oracle_rtvquery	OC-CACHE OC-CLUSTER OC-CLUST
RTV-DATA-SOLMON	✓	56	212	https://rtvdemos.sl.com/solmon_rtvquer	SOLACE-BRIDGE SOLACE-CLIENT S
RTV-DATA-TIBCO	✓	76	1475	https://rtvdemos.sl.com/tibmon_rtvquer	AMX-NODE AMX-SERVICE AMX-SER
Z-SIMDATA-1	✓	28	577	https://rtvdemos-163.sl.com/simdata_rt	BW-ENGINE BW-PROCESS BW-SER

## RTView Data Servers - HTML

Check the connection status of all your data servers, compare their **Receive Count** values, get their **Connection String** and **Config** which is the RTView version installed.

Each row in the table contains data for a particular data server. You can search, filter, sort and choose columns to include by clicking a column header icon (to the right of each column label) and selecting **Filter**, **Sort Ascending**, **Sort Descending** or **Columns**. Right-click on a table cell to **Export to Excel**. Use **Ctrl +** click or **Shift +** click to select multiple alerts.

To investigate a data server, double-click a row to view detailed metrics for the data server in the ["Data Server Summary - HTML"](#) display.

**RTView Data Server Summary** 02-May-2019 16:56 ✓ DATA

Data Server: RTV-DATA-BOOMI

Name: RTV-DATA-BOOMI      Connected: true      Status: OK  
 Receive Count: 1309      Receive Time: 02-May-2019 16:56:41  
 Config: APM.4.2.1.0\_20181016\_001.28982-beta\_314      Connection String: https://rtvdemos.sl.com/boomimon\_rtvdata/

**Caches**

CacheName	Rows	Columns	Memory
BoomiAtom	50	24	25677
BoomiEvent	3	19	4519
BoomiExecution	203	17	43014
BoomiMolecule	25	16	6841
BoomiMoleculeExecution	25	16	6318
BoomiMoleculeIncoming	25	10	4772
BoomiMoleculeOutgoing	25	12	5185
BoomiMoleculeStats	25	16	6624
BoomiNode	153	26	75604
BoomiNodeStats	153	31	109320
BoomiOutgoing	203	13	37038
BoomiRegisteredAtom	0	11	1028

Page 1 of 2      1 - 40 of 46 items

## Data Server Summary - HTML

Investigate cache table sizes and memory utilization on a data server. Answer questions such as: Are the numbers of each type of data as expected for my data server?

Each row in the table contains data for a particular cache for the data server selected from the drop-down menu. You can see the number of **Rows** and **Columns** are in each cache as well as the amount of **Memory** used by each.

You can search, filter, sort and choose columns to include by clicking a column header icon (to the right of each column label) and selecting **Filter**, **Sort Ascending**, **Sort Descending** or **Columns**. Or just click a column header to sort.

Right-click on a table cell to **Export to Excel** or **Copy Cell Value**. Use **Ctrl +** click or **Shift +** click to select multiple alerts.

**RTView Data Servers Table** 02-May-2019 16:55 ✓ DATA

Local Connections to Data Server

Name	Connected	Status	Connection String	Receive Count	Receive Time	
RTV-DATA-INFRA	✓	OK	https://rtvdemos.sl.com/inframom_rtvddata/	2787	02-May-2019 16:55:15	AP
RTV-DATA-KAFKAMON	✓	OK	https://rtvdemos.sl.com/kafkamom_rtvddata/	2259	02-May-2019 16:55:15	AP
RTV-DATA-ORACLE	✓	OK	https://rtvdemos.sl.com/oracle_rtvddata/	7	02-May-2019 12:05:47	AP
RTV-DATA-TIBCO	✓	OK	https://rtvdemos.sl.com/tibcomom_rtvddata/	5739	02-May-2019 16:55:19	AP
RTV-DATA-BOOMI	✓	OK	https://rtvdemos.sl.com/boomimom_rtvddata/	1301	02-May-2019 16:55:18	AP
RTV-DATA-SOLMON	✓	OK	https://rtvdemos.sl.com/solmcomom_rtvddata/	1504	02-May-2019 16:54:59	AP
RTV-DATA-IBM	✓	OK	https://rtvdemos.sl.com/ibmmomom_rtvddata/	517	02-May-2019 16:53:49	AP
Z-SIMDATA-1	✓	OK	https://rtvdemos-163.sl.com/simdata_rtvddata/	392	02-May-2019 16:53:27	AP

## CI Type Definitions - HTML

View all CI Types that are in your system, investigate their definitions and Service Model mapping, find out the name of their associated cache and their associated alert.

The **CI Type Definitions** table contains all CI Types that are defined in your system. Each row of data describes a particular CI Type, including:

- **INDEXMAP:** The number of indexes and the order in which they are used to form the CI Name.
- **INDEXNAMES:** A semicolon-separated list of the index columns.
- **RTVDISPLAY:** The name of the RTView display to drill-down to from the Alerts Table to see summary data for this CI Type. This is the target of the Go To CI button in the Alerts Table and in the Service Summary display.
- **CIVARMAP:** The names of substitutions that must be set to drill-down to the display.
- **DEFAULTQUALITY:** A flag indicating whether the lack of data is considered an error condition or not.
- **OWNER:** The Owner the CIType is associated with, when the CMDB is populated automatically from CIs of this type.
- **AREA:** The Area the CIType is associated with.
- **SERVICEGROUP:** The SERVICEGROUP the CIType is associated with, when the CMDB is populated automatically from CIs of this type.

The **Cache Map By CI Type** table lists all CI Types that are defined in your system and the name of the cache associated with each.

The **Alert Map By CI Type** table lists all CI Types that are defined in your system and the name of the alerts associated with each.

You can search, filter, sort and choose columns to include by clicking a column header icon (to the right of each column label) and selecting **Filter**, **Sort Ascending**, **Sort Descending** or **Columns**. Or just click a column header to sort.

Right-click on a table cell to **Export to Excel** or **Copy Cell Value**. Use **Ctrl +** click or **Shift +** click to select multiple alerts.

02-May-2019 17:00 ✓ DATA

### CI Type Definitions

CITYPE	INDEXMAP	INDEXNAMES	RTVDISPLAY	CIVARMAP	DEFAULTQU...	OWNER	AREA	SERVICEGR...
EM-SERVICE	1,2,3,4,5	Owner:Area:Serv	rtv_service_summ	SrvOwnerTemp;	1	Services	DependentAreas	DependentGroup
ACW	1	Dimension	acw_instance_su	SawsEc2Instance	1	Infrastructure	Servers	Hosts
AMX-SERVICE	1,2	Application:Serv	amx_service_sur	\$amxApplication;	1	Infrastructure	Middleware	TIBCO-AMX
AMX-SERVICEN	1,2,3,4	AMX Host:Node;	amx_servicenode	\$amxHost;\$amxof	1	Infrastructure	Middleware	TIBCO-AMX
AMX-NODE	1,2,3	HostName:AMX	amx_node_summ	\$HostName;\$amx	1	Infrastructure	Middleware	TIBCO-AMX
AMX-HOST	1	AMX Host	amx_host_summ	\$amxHost	1	Infrastructure	Middleware	TIBCO-AMX
BW6-APP	1,2,3	Domain:AppSpac	bw6_app_summ	\$bw6domain;\$bv	1	Infrastructure	Middleware	TIBCO-BW6
BW6-APPNODE	1,2,3	Domain:AppSpac	bw6_appnode_si	\$bw6domain;\$bv	1	Infrastructure	Middleware	TIBCO-BW6

Page 1 of 2 1 - 40 of 76 items

### Cache Map By CI Type

CITYPE	CACHENAME
EM-SERVICE	RtvCmdbServiceTable_local
EM-SERVICE	RtvCmdbServiceStats_local
ACW	AwsEc2InstanceStats
AMX-SERVICE	AmxServiceTotals
AMX-SERVICENODE	AmxServices
AMX-NODE	AmxNodes
BW6-APP	Bw6Apps

Page 1 of 3 1 - 40 of 103 items

### Alert Map By CIType

CITYPE	ALERTNAME
EM-SERVICE	RtvEmServiceAlert
EM-SERVICE	RtvEmServiceAlertImpactHigh
ACW	AcwInstanceCpuHigh
ACW	AcwInstanceDiskReadBytesHigh
ACW	AcwInstanceDiskReadOpsHigh
ACW	AcwInstanceDiskWriteBytesHigh
ACW	AcwInstanceDiskWriteOpsHigh

Page 1 of 13 1 - 40 of 494 items

## CI Stats - HTML

Investigate an alert and the CI associated with the alert. This display provides a list of CIs that currently have an active warning or alarm alert. You can find out where the CIs are (which cache they reside on) and the name of the CI associated with the alert.

The **Alert Stats By CI** table lists all CIs that currently have active warning or alarm alerts and provides the **CITYPE** for each.

The **Cache Map By CIType** table lists all **CIType**s for which we are receiving data, provides the associated **CACHENAME** and the associated **Source** RTView DataServer.

The **Cache Map By CI** table lists all **CI**s that we are receiving from the data servers, and provides the associated **CIType**.

You can search, filter, sort and choose columns to include by clicking a column header icon (to the right of each column label) and selecting **Filter**, **Sort Ascending**, **Sort Descending** or **Columns**. Or just click a column header to sort.

Right-click on a table cell to **Export to Excel** or **Copy Cell Value**. Use **Ctrl +** click or **Shift +** click to select multiple alerts.

**CI Stats Tables** 02-May-2019 17:04 ✓ DATA

**Alert Stats By CI**

CITYPE	CINAME
SOLACE-MSGROUTER	VMR-47
KAFKA-CLUSTER	KafkaCRM
BW-SERVER	SLOL6-1(slappm)
BW-SERVER	SLHOST16(sl_qa_conn)
BW-ENGINE	slappm(slappm);domainslappm.BWEngine.Process Space Archive
BW-ENGINE	slappm(slappm);domainslappm.BWApp-5.Procs
BW-ENGINE	slel4-84(slmon);domainslmon.BWApp-4.Procs

Page 1 of 2 1 - 40 of 57 items

**Cache Map By CI Type**

CITYPE	CACHENAME	Source
SOLACE-MSGROUTER	SolAppliances	RTV-DATA-SOLMON
SOLACE-VPN	SolVpns	RTV-DATA-SOLMON
SOLACE-BRIDGE	SolBridges	RTV-DATA-SOLMON
SOLACE-CLIENT	SolClients	RTV-DATA-SOLMON
SOLACE-ENDPOINT	SolEndpoints	RTV-DATA-SOLMON
JVM	JvmConnections	RTV-DATA-KAFKAMON
.JVM	.JvmOperationSystem	RTV-DATA-KAFKAMON

Page 1 of 3 1 - 40 of 115 items

**Cache Map By CI**

CITYPE	CINAME
JVM	localhost:SOLMON_TOMCAT
JVM	localhost:EMSMON_HISTORIAN
JVM	localhost:EMSMON_DATASERVER
JVM	localhost:EMSMON_TOMCAT
JVM	localhost:EMSMON_DATABASE
JVM	localhost:SOLMON_DATASERVER
.JVM	localhost:SOLMON_DISPLAYSERVER

Page 1 of 71 1 - 40 of 2803 items

## Cache Table - HTML

View the raw data that RTView is capturing and maintaining to investigate utilization and capacity metrics, as well as connection details, for caches on a data server.

Select a **Data Server** from the drop-down menu. The upper table contains a row of data for each cache on the selected data server. You can see the current number of **Rows** and **Columns** in each table and the amount of **Memory** used. You can also find out the cache **Table** type of which there are five:

- **current** tables show the most recently received values for each index.
- **current\_condensed** tables are current tables with primary compaction configured.
- **history** tables show the historical values for each index.
- **history\_condensed** tables are history tables with primary compaction configured.
- **history\_combo** tables are history tables with primary compaction configured, and which is also configured to store rows of recent raw data followed by rows of older condensed data.

Select a cache to see connection utilization details for that cache in the lower table. The lower table shows the contents of the selected cache table. Available columns vary by cache. For example, a JVM cache table might provide **BootClassPath** and **InputArgument** columns, and a Tomcat cache might provide **RateAccess** and **cacheMaxSize** columns.

You can search, filter, sort and choose columns to include by clicking a column header icon (to the right of each column label) and selecting **Filter**, **Sort Ascending**, **Sort Descending** or **Columns**. Or just click a column header to sort.

Right-click on a table cell to **Export to Excel** or **Copy Cell Value**. Use **Ctrl + click** or **Shift + click** to select multiple alerts. Use **History Tables** to include / exclude history tables in the table. Right-click on a table cell to **Export to Excel** or **Copy Cell Value**.

This low-level option can be useful to identify the source of the problem when the displays are not showing the expected data. Use this data for debugging and troubleshooting with Technical Support.

**Cache Table** 07-May-2019 14:11 ✓ DATA

Data Server: central-alert History Tables:

Data Server URL: [https://rtvdemos.sl.com/emdemo\\_central\\_rtvquery](https://rtvdemos.sl.com/emdemo_central_rtvquery)

Cache	Table	Rows	Columns	Memory
JmxStatsTotals	current	1	4	441
RtvAlertGroupMap	current	493	3	67424
RtvAlertMapByCI	current	62	5	13614
RtvAlertSourceStats	current	8	2	940
RtvAlertStatsByArea	current	8	9	2930
RtvAlertStatsByAreaAndAlertGroup	current	8	10	3454
RtvAlertStatsByCI	current	59	5	9228
RtvAlertStatsByCIAndAlertGroup	current	59	6	12506

Cache: RtvAlertStatsByCIAndAlertGroup Table: current

time_stamp	CITYPE	CINAME	ALERTGROUP	MaxSeverity	AlertCount
2019-May-07 14:11:33	JVM	localhost:SOLOMON_TOM	None	2	1
2019-May-07 14:11:33	JVM	localhost:EMSMON_DAT	None	2	1
2019-May-07 14:11:33	JVM	localhost:SOLMON_DISF	None	2	1
2019-May-07 14:11:33	JVM	localhost:SOLMON_DAT	None	2	1
2019-May-07 14:11:33	JVM	localhost:EMSMON_DISI	None	2	1
2019-May-07 14:11:33	JVM	localhost:SOLMON_TOM	None	2	1
2019-May-07 14:11:33	JVM	localhost:EMSMON_DAT	None	2	1
2019-May-07 14:11:33	JVM	Instance-1-90;CRMBroke	None	1	1
2019-May-07 14:11:33	JVM	Instance-1-90;CRMZooki	None	1	1
2019-May-07 14:11:33	JVM	Instance-1-171;CRMCon	None	1	1
2019-May-07 14:11:33	JVM	Instance-1-171;CRMCon	None	1	1
2019-May-07 14:11:33	JVM	Instance-1-171;CRMBrok	None	1	1
2019-May-07 14:11:33	JVM	localhost:TMolecule5_2	None	1	1
2019-May-07 14:11:33	JVM	localhost:PMolecule12_1	None	1	1

Page 1 of 2 1 - 40 of 59 items



## CHAPTER 3 RTView DataServer for IBM

The RTView DataServer for IBM provides a way to create connections and modify default configuration settings for solution packages and sends collected data to RTViewCentral. RTViewCentral contains the displays associated with the RTView DataServer for IBM which you use to monitor your IBM components. Both the display server user interface and the HTML user interface are described here.

The RTView *DataCollector* for IBM is also available for use with the RTView DataServer for IBM. RTView DataCollector for IBM is used for collecting and sending data to one or more data servers. The RTView DataCollector for IBM is also useful if you need to distribute data collection.

For an overview and details about configuring RTView Enterprise, including RTViewCentral, RTView DataServers, RTView DataCollectors and solution packages, see the *RTView Enterprise Configuration Guide*.

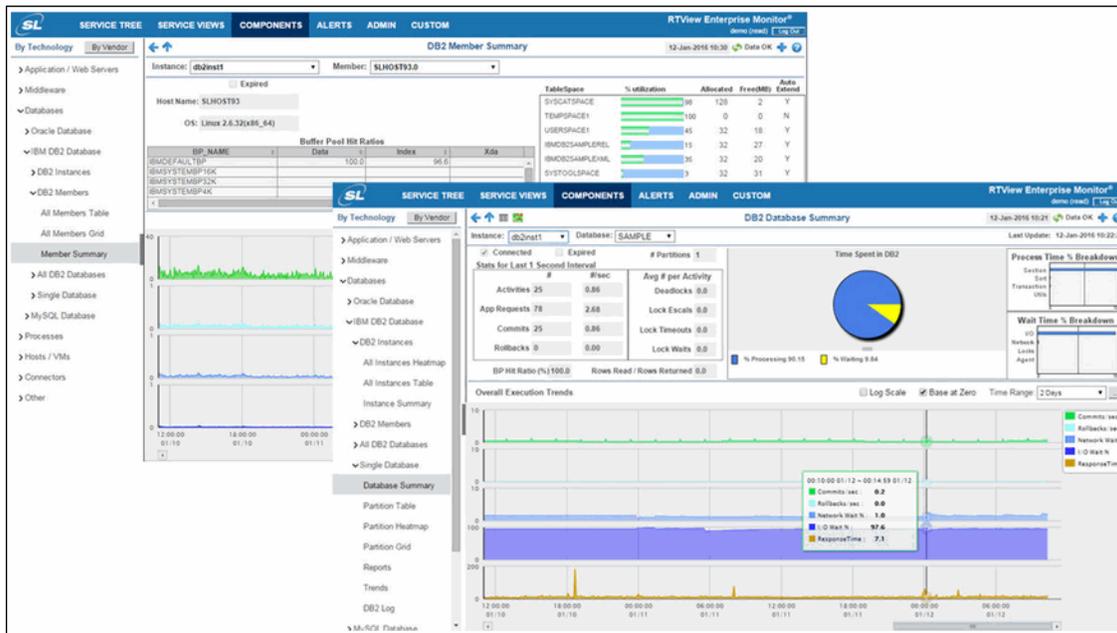
RTViewCentral contains the following solution packages and displays that will be populated with data collected via the RTView DataServer for IBM:

- "IBM DB2"
- "IBM DB2 - HTML"
- "IBM MQ"
- "IBM MQ - HTML"
- "IBM WebSphere"
- "IBM WebSphere - HTML"

**Note:** This document assumes familiarity with the products monitored. For additional details, refer to vendor documentation.

## IBM DB2

The Solution Package for IBM® DB2 includes a high level heatmap and tabular displays as well as drilldown views to access real-time and historical performance metrics for each DB2 Database in your monitored services and applications.



With the Solution Package for IBM DB2, you are able to drill down from a high level alert at a business service or application health level into the supporting database infrastructure, to determine what is causing the alert and to take corrective action. This service-centric approach makes it easy for application support teams and IBM DBAs to prioritize incidents based on the impact to the business.

Solution packages include a data adapter, real-time memory cache, alert rule engine, pre-configured displays, and a data historian for persisting of real-time performance metrics.

The following IBM DB2 Views can be found under **Components** tab > **Databases** > **IBM DB2 Database**:

- "DB2 Instances"
- "DB2 Members"
- "All DB2 Databases"
- "Single Database"

## DB2 Instances

Displays in this View are:

- "All Instances Heatmap"
- "All Instances Table"
- "Instance Summary"

## All Instances Heatmap

View current alert status and performance metrics of all or just one of your IBM DB2 instances. Use the **Metric** drop-down menu to view **Alert Severity, Alert Count, Response Time, I/O Wait Time %, Network Wait Time %, Agent Wait Time %, Avg Deadlocks per Activity, Avg Lock Escalations per Activity, Avg Lock Timeouts per Activity, Avg Lock Waits per Activity, Rows Read per Rows Returned, Activites/sec, App Requests/sec, App Commits/sec, App Rollbacks/sec** or **Buffer Pool Hit Ratio %**.

Answer questions such as, Are any instances reaching a state of critical health? Do I need to allocate more tablespace to an instance? Are response times slow on any instances? Are application deadlocks causing bottlenecks on any instances? Is processing load evenly distributed across instances?

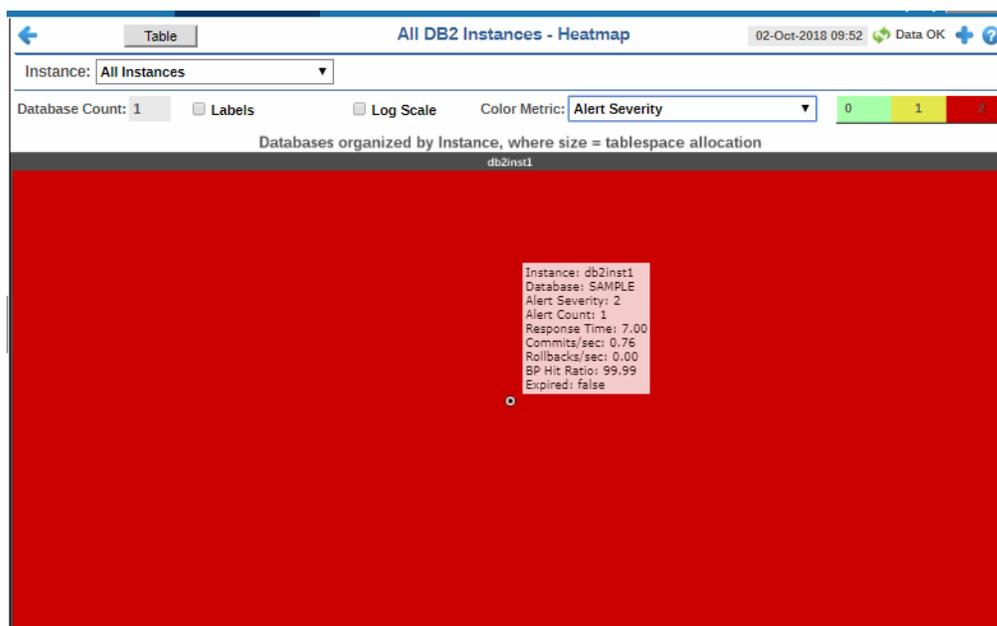
Each rectangle in the heatmap represents a different instance, where the rectangle color indicates the most critical alert state for items associated with that instance, and the rectangle size represents the tablespace allocation size for the instance.

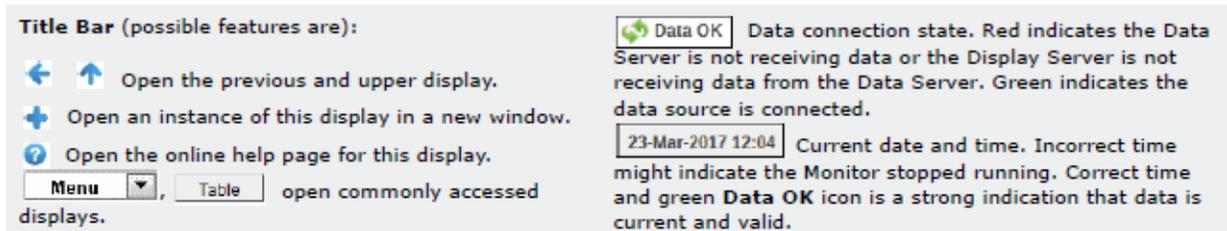
Each metric has its own color gradient bar legend that maps values to colors. By default, the Alert Severity metric is shown, which is the current alert severity for items associated with the rectangle. Values range from **0** - **2**, as indicated in the color gradient  bar:

- **(2)** Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
- **(1)** Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
- **(0)** Green indicates that no metrics have exceeded their alert thresholds.

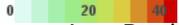
Other performance metrics include wait times, application requests and rollbacks.

Use the **Labels** check-box  to include or exclude labels in the heatmap, use the **Log Scale** check-box  to apply log scale and mouse over a rectangle to see additional metrics. Click a rectangle to drill down on an instance in the **"Instance Summary"** display.





## Fields and Data:

- Instance:** Select an instance.
- Labels** Select this check box to display the names of the instances at the top of each rectangle in the heatmap.
- Log Scale** Select this check box to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Metric** Choose a metric to view in the display. For details about the data, refer to vendor documentation.
- Alert Severity** The current alert severity for items associated with the rectangle. Values range from **0** - **2**, as indicated in the color gradient  bar, where **2** is the highest Alert Severity:
- (2)** Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
  - (1)** Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
  - (0)** Green indicates that no metrics have exceeded their alert thresholds.
- Alert Count** The total number of critical and warning unacknowledged alerts for items associated with the rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
- Response Time** The average response time, in milliseconds, for items associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of the **Db2ResponseTimeHigh** alert. The middle value in the gradient bar indicates the middle value of the range.
- (2)** Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
  - (1)** Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
  - (0)** Green indicates that no metrics have exceeded their alert thresholds.

**I/O Wait Time %**

The percentage of the wait time being used by I/O processes. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0%** to **100%** of the **Db2ResponseTimeHigh** alert threshold. The middle value in the gradient bar indicates the middle value of the range.

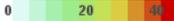
 Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.

 Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.

 Green indicates that no metrics have exceeded their alert thresholds.

The percentage wait time taken by I/O operations.

**Network Wait Time %**

The percentage of the wait time being used by network processes. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0%** to **100%** of the **Db2ResponseTimeHigh** alert threshold. The middle value in the gradient bar indicates the middle value of the range.

 Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.

 Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.

 Green indicates that no metrics have exceeded their alert thresholds.

**Agent Wait Time %**

The percentage of the wait time being used by agent processes. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0%** to **100%** of the **Db2ResponseTimeHigh** alert threshold. The middle value in the gradient bar indicates the middle value of the range.

 Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.

 Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.

 Green indicates that no metrics have exceeded their alert thresholds.

**Avg Deadlocks per Activity**

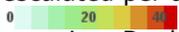
The average number of application deadlocks per activity. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of the **Db2DeadlockRateHigh** alert. The middle value in the gradient bar indicates the middle value of the range.

 Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.

 Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.

 Green indicates that no metrics have exceeded their alert thresholds.

**Avg Lock Escalations per Activity**

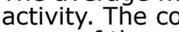
The average number of application deadlocks that were escalated per activity. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of the **Db2DeadlockRateHigh** alert. The middle value in the gradient bar indicates the middle value of the range.

 Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.

 Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.

 Green indicates that no metrics have exceeded their alert thresholds.

**Avg Lock Timeouts per Activity**

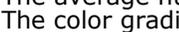
The average number of application lock timeouts per activity. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of the **Db2LockWaitTimeHigh** alert. The middle value in the gradient bar indicates the middle value of the range.

 Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.

 Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.

 Green indicates that no metrics have exceeded their alert thresholds.

**Avg Lock Waits per Activity**

The average number of application lock waits per activity. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of the **Db2LockWaitTimeHigh** alert. The middle value in the gradient bar indicates the middle value of the range.

 Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.

 Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.

 Green indicates that no metrics have exceeded their alert thresholds. The average number of application deadlock timeouts per activity.

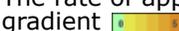
**Rows Read per Rows Returned**

The total number of rows read per number of rows returned. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum value in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

**Activites/sec**

The rate of activities (per second). The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum number of incoming messages per second. The middle value in the gradient bar indicates the middle value of the range.

**App Requests/sec**

The rate of application requests (per second). The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum number of incoming messages per second. The middle value in the gradient bar indicates the middle value of the range.

**App Commits/sec**

The rate of application commits (per second). The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum number of incoming messages per second. The middle value in the gradient bar indicates the middle value of the range.

**App Rollbacks/sec**

The rate of application rollbacks (per second). The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum number of incoming messages per second. The middle value in the gradient bar indicates the middle value of the range.

**Buffer Pool Hit Ratio %**

The percentage Tablespace used by the buffer pool hit ratio, which is the total number of pool hits divided by the total number of buffer pool lookups. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0%** to **100%** of the **Db2UsedTablespaceHig** alert threshold. The middle value in the gradient bar indicates the middle value of the range.

 Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.

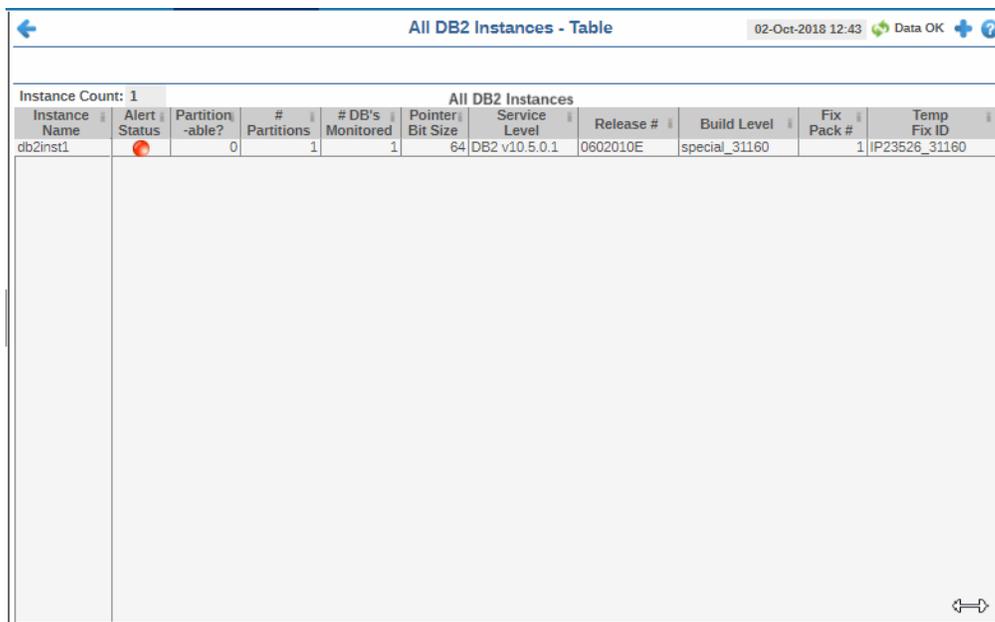
 Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.

 Green indicates that no metrics have exceeded their alert thresholds.

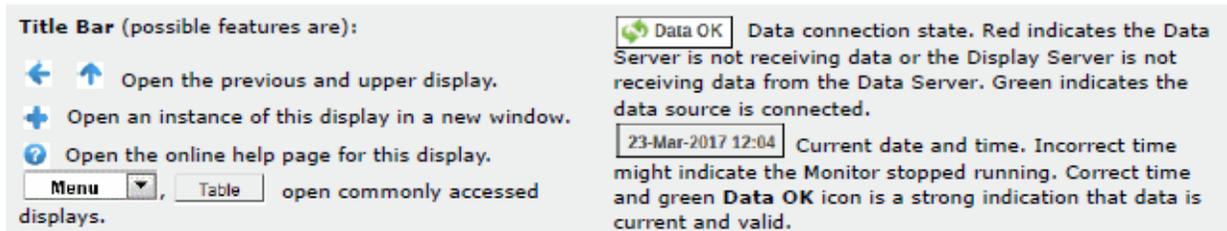
**All Instances Table**

This display provides a list of all DB2 instances, configuration details about each instance and their alert status. Each row in the table is a different DB2 instance.

Click a column header to sort column data in numerical or alphabetical order. Investigate by clicking a row to view details for an instance in the "Instance Summary" display.



All DB2 Instances - Table										
Instance Count: 1										
All DB2 Instances										
Instance Name	Alert Status	Partitionable?	# Partitions	# DB's Monitored	Pointer Bit Size	Service Level	Release #	Build Level	Fix Pack #	Temp Fix ID
db2inst1		0	1	1	64	DB2 v10.5.0.1	0602010E	special_31160	1	IP23526_31160



<b>Instance Count</b>	The number of instances in the table.
<b>Instance Name</b>	The name of the instance.
<b>Alert Status</b>	<ul style="list-style-type: none"> <li>● Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li>● Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li>● Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Partitionable?</b>	Indicates whether the instance can be partitioned. <b>0</b> is No <b>1</b> is Yes
<b># Partitions</b>	The number of partitions for the instance.
<b>#DB's Monitored</b>	The number of databases monitored.
<b>Pointer Bit Size</b>	The pointer bit size (e.g. 64).
<b>Service Level</b>	The installed DB2 software release version.
<b>Release #</b>	The installed DB2 software release number.
<b>Build Level</b>	The installed DB2 software build number.
<b>Fix Pack #</b>	The number of the installed DB2 Fix Pack.
<b>Temp Fix ID</b>	The IBM fix ID.

## Instance Summary

Select an instance **Instance** to see the following information for that instance:

The **Instance Members** table shows resource allocation and configuration details (CPU Load, Memory) for members (hosts) on a single IBM DB2 instance. Click a row to drill-down to details in the **“Member Summary”** display.

The **Instance Databases** table shows utilization and processing metrics for databases on a single IBM DB2 instance. Click a row to drill-down to details in the **“Database Summary”** display.

The screenshot shows the 'DB2 Instance Summary' window for instance 'db2inst1'. It features a title bar with navigation icons, a 'Data OK' status indicator, and a timestamp '02-Oct-2018 13:10'. Below the instance name, there are two main sections:

**Instance Members** (Member Count: 1)

Host Name	Expired	Member Count	OS	CPU Count	CPU Speed	Total RA	Total Virtu. Memory	CPU Lo. Short	CPU Lo. Mediu...	CPU Lo. Long	% CPU Usage
SLHOST93	<input type="checkbox"/>	0	Linux 2.6.32(x86_64)	1 of 1 in use	2,300	3,835	7,867	0.26	0.14	0.10	26.00

**Instance Databases** (Database Count: 1)

Database	Expired	Connect.	Severi..	Respons.. Time (ms...	Activiti.. / sec	App Requests / sec	Commi.. / sec	Rollbacks / sec	Activity W.. Time %	Agent Wait Time %	IO W... Time ...	Network Time %
SAMPLE	<input type="checkbox"/>	<input checked="" type="checkbox"/>		4.00	0.89	2.78	0.89	0.00	10.45	0.00	98.50	

**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Instance Members Table**

Each row is a different host member. Column values describe the host.

- Member Count:** The number of members in the table.
- Host Name** The name of the host.
- Expired** When checked, performance data has not been received within the time specified by your administrator for the **Expire Time**. If your administrator has also set the **Delete Time**, this row will be deleted if no data is received within the time specified for deletion.
- Member** The member number.
- OS** The installed operating system.

<b>CPUs</b>	The number of CPUs and the number of CPUs in use.
<b>CPU Speed</b>	The processor speed.
<b>Total RAM</b>	The total amount of RAM, in megabytes.
<b>Total Virtual Memory</b>	The total amount of virtual memory, in megabytes.
<b>CPU Load Short</b>	Amount of processor load over the short term (defined by the IBM DB2 system, for example, 1-5 minutes).
<b>CPU Load Medium</b>	Amount of processor load over the medium term (defined by the IBM DB2 system, for example, 5-10 minutes).
<b>CPU Load Long</b>	Percentage of CPU load over the long term (defined by the IBM DB2 system, for example, 10-15 minutes).
<b>% CPU Usage</b>	The percentage of CPU used.

**Instance Databases Table**

Each row is a different database. Column values describe the database.

**Database Count:** The number of databases in the table.

<b>Database</b>	The name of the database.
<b>Expired</b>	When checked, performance data has not been received within the time specified by your administrator for the <b>Expire Time</b> . If your administrator has also set the <b>Delete Time</b> , this row will be deleted if no data is received within the time specified for deletion.
<b>Connected</b>	When checked, the database is connected.
<b>Severity</b>	The alert status:  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
<b>Response Time</b>	The response time, in milliseconds.
<b>I/O Wait Time %</b>	The percentage wait time taken by I/O operations.
<b>Network Wait Time %</b>	The percentage wait time taken by the network.
<b>Agent Wait Time %</b>	The percentage wait time taken by agents.
<b>Avg Deadlocks per Activity</b>	The average number of application deadlocks per activity.
<b>Avg Lock Escalations per Activity</b>	The average number of application deadlock escalations per activity.
<b>Avg Lock Timeouts per Activity</b>	The average number of application deadlock timeouts per activity.
<b>Avg Lock Waits per Activity</b>	The average number of application deadlock waits per activity.
<b>Rows Read per Rows Returned</b>	The number of rows read per number of rows returned.
<b>Activites/sec</b>	The number of activities per second.

<b>App Requests/sec</b>	The number of application requests per second.
<b>Commits/sec</b>	The number of application commits per second.
<b>Rollbacks/sec</b>	The number of application rollbacks per second.
<b>Buffer Pool Hit Ratio %</b>	The current buffer pool hit ratio, which is the total number of pool hits divided by the total number of buffer pool lookups.
<b>Activity Wait Time %</b>	The percentage wait time taken by activities.
<b>Avg Request CPU Time</b>	The average amount of CPU time used by requests, in seconds.
<b>Compile Proc Time %</b>	The percentage of time used for compiling processes.
<b>Routine Time Request %</b>	The percentage of time used for routine request processes.
<b>Section Time %</b>	The percentage of time used for section processes.
<b>Section Sort Time %</b>	The percentage of time used for sorting section processes.
<b>BP Hit Ratio %</b>	The current buffer pool hit ratio, which is the total number of pool hits divided by the total number of buffer pool lookups.
<b>Transaction Time %</b>	The percentage of time used for transaction processes.
<b>Utils Proc Time %</b>	The percentage of time used for utilities processes.
<b>Timestamp</b>	The data and time this data was last updated.

## DB2 Members

Displays in this View are:

- ["All Members Table"](#): Get configuration details and utilization metrics for host members of an instance.
- ["All Members Grid"](#): View trend graph for CPU load and CPU utilization for each member of an instance.
- ["Member Summary"](#): Investigate TableSpaces and buffer pool utilization metrics for a host member.

### All Members Table

Select an instance **Instance** to see a list of all host members of an instance.

Each row is a different host that shows configuration details and utilization metrics for each. Details include OS, RAM, CPU load (short, medium, long) and virtual memory.

Click a column header to sort column data in numerical or alphabetical order. Drill-down and investigate by clicking a row to view details for a host member in the “Member Summary” display.

Host Name	Expir..	Instance	Mem..	OS	CPUs	CPU Speed	Total R...	Total Virt... Memor...	CPU L... Shor...	CPU Lo... Mediu...	CPU L... Lon...	% Us
SLHOST93	<input type="checkbox"/>	db2inst1	0	Linux 2.6.32(x86_64)	1 of 1 in use	2,300	3,835	7,867	0.12	0.07	0.02	

#### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- ,  open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

### Instance Members Table

Each row is a different host member. Column values describe the host.

**Member Count:** The number of members in the table.

**Host Name** The name of the host.

**Expired** When checked, performance data has not been received within the time specified by your administrator for the **Expire Time**. If your administrator has also set the **Delete Time**, this row will be deleted if no data is received within the time specified for deletion.

**Instance** The name of the instance that the host is a member of.

**Member** The name of the member.

**OS** The installed operating system.

**CPUs** The number of CPUs and the number of CPUs in use.

**CPU Speed** The processor speed.

**Total RAM** The total amount of RAM, in megabytes.

<b>Total Virtual Memory</b>	The total amount of virtual memory, in megabytes.
<b>CPU Load Short</b>	Amount of processor load over the short term (defined by the IBM DB2 system, for example, 1-5 minutes).
<b>CPU Load Medium</b>	Amount of processor load over the medium term (defined by the IBM DB2 system, for example, 5-10 minutes).
<b>CPU Load Long</b>	Percentage of CPU load over the long term (defined by the IBM DB2 system, for example, 10-15 minutes).
<b>% CPU Usage</b>	The percentage of CPU used.
<b>Timestamp</b>	The data and time this data was last updated.

## All Members Grid

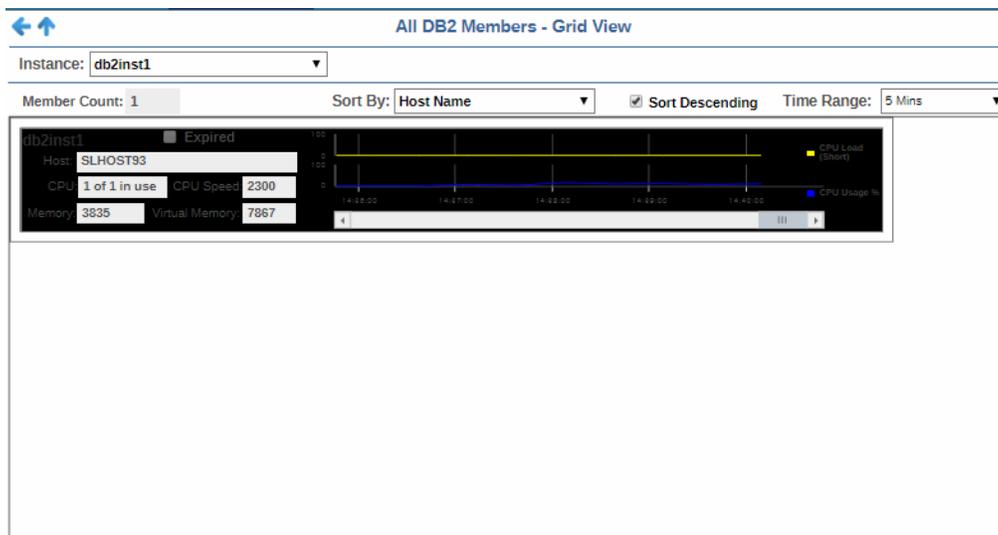
This display provides a grid view of a subset of data in the ["All Members Table"](#) display.

Each grid object is a different host member of the selected instance. Each has a trend graph that traces the **CPU Load Short** and **% CPU Usage** for the host member. Choose a **Time Range** for the trend graph to trace, or select **All Data** to include all available data in the trace. Scroll forward and backward in the trend graph and click **Reset** to return to the original state.

Show **Expired** members first using **Sort By: Expired**. Or show members with most severe **CPU Load Short** values first, and most severe **CPU Usage** values first.

Toggle **Sort Descending** to order the grid objects.

Investigate performance metrics details for a host member in the ["Member Summary"](#) display by clicking on the (left or right side of the) grid object.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

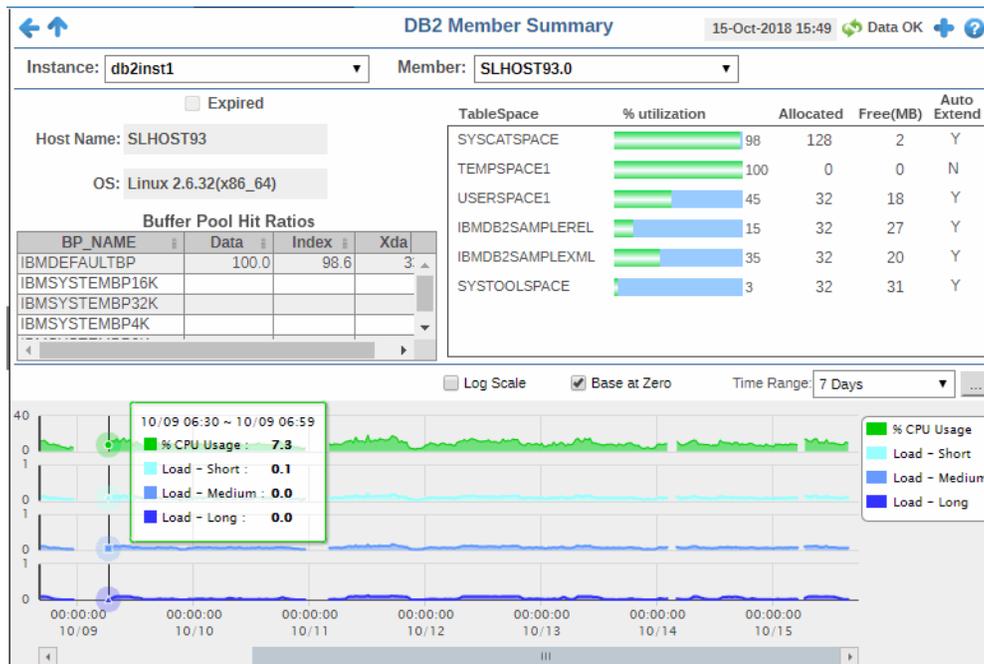
<b>Instance:</b>	Select an instance.														
<b>Member Count:</b>	The number of members in the display.														
<b>Sort By:</b>	Orders the grid objects as follows: <ul style="list-style-type: none"> <li>• <b>Instance:</b> Instance names.</li> <li>• <b>Host Name:</b> Host member names.</li> <li>• <b>CPU Load - Short:</b> Current host member <b>CPU Load - Short</b> values, from most to least critical.</li> <li>• <b>CPU Usage:</b> Current host member <b>CPU Usage</b> values, from most to least critical.</li> <li>• <b>Expired:</b> Host members that are in an expired state.</li> </ul>														
<b>Sort Descending</b>	Toggle on to order the grid objects.														
<b>Time Range</b>	Choose a time range for the trend graph to trace, from 2 minutes to 7 days, or choose <b>All Data</b> to include all available data. The selected time range applies to all grid objects in the display.														
<b>Grid Objects</b>	Each grid object is a different host member and values describe each host member. <table> <tr> <td><b>Expired</b></td> <td>When checked, performance data has not been received within the time specified by your administrator for the <b>Expire Time</b>. If your administrator has also set the <b>Delete Time</b>, this grid object will be deleted if no data is received within the time specified for deletion.</td> </tr> <tr> <td><b>Host</b></td> <td>The name of the host.</td> </tr> <tr> <td><b>CPU</b></td> <td>The number of CPUs and the number of CPUs in use.</td> </tr> <tr> <td><b>CPU Speed</b></td> <td>The processor speed.</td> </tr> <tr> <td><b>Memory</b></td> <td>The total amount of memory, in megabytes.</td> </tr> <tr> <td><b>Virtual Memory</b></td> <td>The total amount of virtual memory, in megabytes.</td> </tr> <tr> <td><b>Trend Graph</b></td> <td> <ul style="list-style-type: none"> <li>• <b>CPU Load Short</b> Traces the amount of CPU load over the short term (defined by the IBM DB2 system, for example, 1-5 minutes).</li> <li>• <b>CPU Usage %</b> Traces the percentage of CPU used.</li> </ul> </td> </tr> </table>	<b>Expired</b>	When checked, performance data has not been received within the time specified by your administrator for the <b>Expire Time</b> . If your administrator has also set the <b>Delete Time</b> , this grid object will be deleted if no data is received within the time specified for deletion.	<b>Host</b>	The name of the host.	<b>CPU</b>	The number of CPUs and the number of CPUs in use.	<b>CPU Speed</b>	The processor speed.	<b>Memory</b>	The total amount of memory, in megabytes.	<b>Virtual Memory</b>	The total amount of virtual memory, in megabytes.	<b>Trend Graph</b>	<ul style="list-style-type: none"> <li>• <b>CPU Load Short</b> Traces the amount of CPU load over the short term (defined by the IBM DB2 system, for example, 1-5 minutes).</li> <li>• <b>CPU Usage %</b> Traces the percentage of CPU used.</li> </ul>
<b>Expired</b>	When checked, performance data has not been received within the time specified by your administrator for the <b>Expire Time</b> . If your administrator has also set the <b>Delete Time</b> , this grid object will be deleted if no data is received within the time specified for deletion.														
<b>Host</b>	The name of the host.														
<b>CPU</b>	The number of CPUs and the number of CPUs in use.														
<b>CPU Speed</b>	The processor speed.														
<b>Memory</b>	The total amount of memory, in megabytes.														
<b>Virtual Memory</b>	The total amount of virtual memory, in megabytes.														
<b>Trend Graph</b>	<ul style="list-style-type: none"> <li>• <b>CPU Load Short</b> Traces the amount of CPU load over the short term (defined by the IBM DB2 system, for example, 1-5 minutes).</li> <li>• <b>CPU Usage %</b> Traces the percentage of CPU used.</li> </ul>														

## Member Summary

Use this display to study trends for a single host member, such as CPU usage and loads (Short, Medium and Long).

Select an **Instance** and a **Member** from the drop-down menus. Check status of the selected host member **TableSpace** allocations and current **Buffer Pool Hit Ratios**.

Choose a **Time Range** or click  to specify your own. Mouse-over to see additional details..



**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu  Table open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Instance:** Choose an instance.

**Member:** Choose a member.

**Expired** When checked, performance data has not been received within the time specified by your administrator for the **Expire Time**. If your administrator has also set the **Delete Time**, this row will be deleted if no data is received within the time specified for deletion.

**Host Name** The name of the selected host member.

**OS** The operating system on the selected host member.

**Buffer Pool Hit Ratios** Each row in the table in a different buffer pool. Values describe values for the selected host member.

**BP\_NAME** The name of the buffer pool.

**Data**

**Index**

**Xda**

### TableSpace

Values describe TableSpace values for the selected host member.

<b>TableSpace</b>	The name of the TableSpace.
<b>% Utilization</b>	The percentage of TableSpace used.
<b>Allocated</b>	The amount of TableSpace allocated.
<b>Free(MB)</b>	The amount of free TableSpace, in megabytes.
<b>Auto Extend</b>	Indicates whether auto extend is enabled. Y/N

### Trend Graph

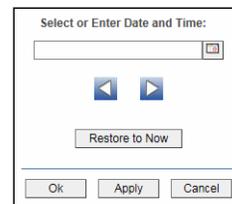
For the selected host member, trend graphs trace as follows:

- **CPU Load Short** Traces the amount of CPU load over the short term (defined by the IBM DB2 system, for example, 1-5 minutes).
- **CPU Load Medium** Traces the amount of CPU load over the medium term (defined by the IBM DB2 system, for example, 5-10 minutes).
- **CPU Load Long** Traces the amount of CPU load over the long term (defined by the IBM DB2 system, for example, 10-15 minutes).
- **CPU Usage %** Traces the percentage of CPU used.

**Log Scale** Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.

**Base at Zero** Select to use zero (**0**) as the Y axis minimum for all graph traces.

**Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. By default, the time range end point is the current time. To enter a specific time range, click the associated ellipsis button .



To change the time range click the Open Calendar button , choose the date and time, then click **OK**. Or, enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss** (for example, Aug 21, 2011 12:24 PM) and click **Apply**. Use the Navigation Arrows   to move forward or backward one time period (the time period selected from the Time Range drop-down menu). Click **Restore to Now** to reset the time range end point to the current time.

## All DB2 Databases

Displays in this View are:

- ["All Databases Table"](#): View a tabular list of performance metrics for all databases in an instance.
- ["All Databases Grid"](#): View trend graphs of databases Commits/sec, Rollbacks/sec and Response Times.

## All Databases Table

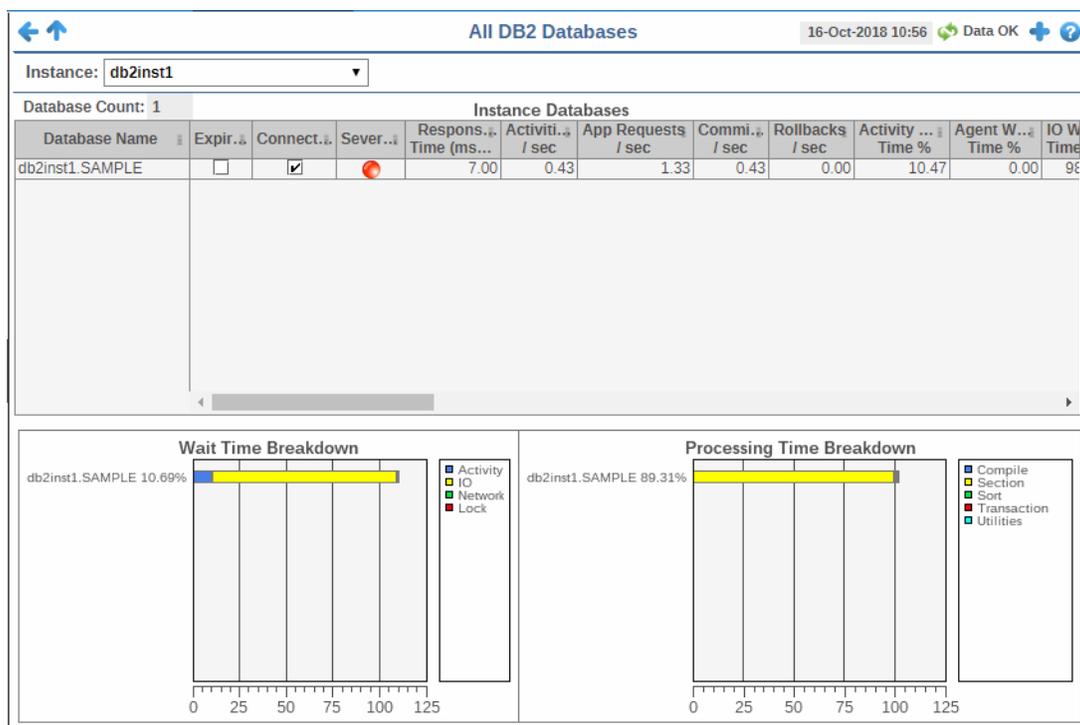
Select an **Instance** and view a list of all its databases and their performance metrics.

Use this display to identify which databases are having or causing issues for the instance.

Sort the list of databases by **Alert Severity** or by any other table column value.

In the bar graphs, view **Wait Time Breakdown** for the instance (by Activity, IO, Network and Lock), as well as **Processing Time Breakdown** (by Compile, Section, Sort, Transaction and Utilities) for the instance.

Click a row to drill-down to details in the ["Database Summary"](#) display.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Instance Databases Table**

Each row is a different database on the selected instance. Column values describe the database.

**Database Count:** The number of databases in the table.

<b>Database</b>	The name of the database.
<b>Expired</b>	When checked, performance data has not been received within the time specified by your administrator for the <b>Expire Time</b> . If your administrator has also set the <b>Delete Time</b> , this row will be deleted if no data is received within the time specified for deletion.
<b>Connected</b>	When checked, the database is connected.
<b>Severity</b>	The alert status: <ul style="list-style-type: none"> <li> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Response Time</b>	The response time, in milliseconds.
<b>I/O Wait Time %</b>	The percentage wait time taken by I/O operations.
<b>Network Wait Time %</b>	The percentage wait time taken by the network.
<b>Agent Wait Time %</b>	The percentage wait time taken by agents.
<b>Avg Deadlocks per Activity</b>	The average number of application deadlocks per activity.
<b>Avg Lock Escalations per Activity</b>	The average number of application deadlock escalations per activity.
<b>Avg Lock Timeouts per Activity</b>	The average number of application deadlock timeouts per activity.
<b>Avg Lock Waits per Activity</b>	The average number of application deadlock waits per activity.
<b>Rows Read per Rows Returned</b>	The number of rows read per number of rows returned.
<b>Activites/sec</b>	The number of activities per second.
<b>App Requests/sec</b>	The number of application requests per second.
<b>Commits/sec</b>	The number of application commits per second.
<b>Rollbacks/sec</b>	The number of application rollbacks per second.
<b>Buffer Pool Hit Ratio %</b>	The current buffer pool hit ratio, which is the total number of pool hits divided by the total number of buffer pool lookups.
<b>Activity Wait Time %</b>	The percentage wait time taken by activities.
<b>Avg Request CPU Time</b>	The average amount of CPU time used by requests, in seconds.
<b>Compile Proc Time %</b>	The percentage of time used for compiling processes.
<b>Routine Time Request %</b>	The percentage of time used for routine request processes.

<b>Section Time %</b>	The percentage of time used for section processes.
<b>Section Sort Time %</b>	The percentage of time used for sorting section processes.
<b>BP Hit Ratio %</b>	The current buffer pool hit ratio, which is the total number of pool hits divided by the total number of buffer pool lookups.
<b>Transaction Time %</b>	The percentage of time used for transaction processes.
<b>Utils Proc Time %</b>	The percentage of time used for utilities processes.
<b>Activity Time (ms)</b>	The amount of time, in milliseconds, used for activity processes.
<b>Agent Idle Time (ms)</b>	The amount of time, in milliseconds, that agents were in an idle state.
<b>CPU Time (ms)</b>	The amount of time, in milliseconds, used for CPU processes.
<b>Processing Time (ms)</b>	The amount of time, in milliseconds, used for processing requests.
<b>Request Time (ms)</b>	The amount of time, in milliseconds, used for processing requests.
<b>DB Processing %</b>	The percentage of time used for database processes.
<b>DB Wait %</b>	The percentage of time used for database waits.
<b># Connections</b>	The current number of database connections.
<b>Timestamp</b>	The data and time this data was last updated.
<b>Wait Time Breakdown</b>	Shows the percentage of total wait time used for the instance categorized by: <ul style="list-style-type: none"> <li>• Activity</li> <li>• IO</li> <li>• Network</li> <li>• Lock</li> </ul>
<b>Processing Time Breakdown</b>	Shows the percentage of total processing time for the instance used by the following types of actions: <ul style="list-style-type: none"> <li>• Compile</li> <li>• Section</li> <li>• Sort</li> <li>• Transaction</li> <li>• Utilities</li> </ul>

## All Databases Grid

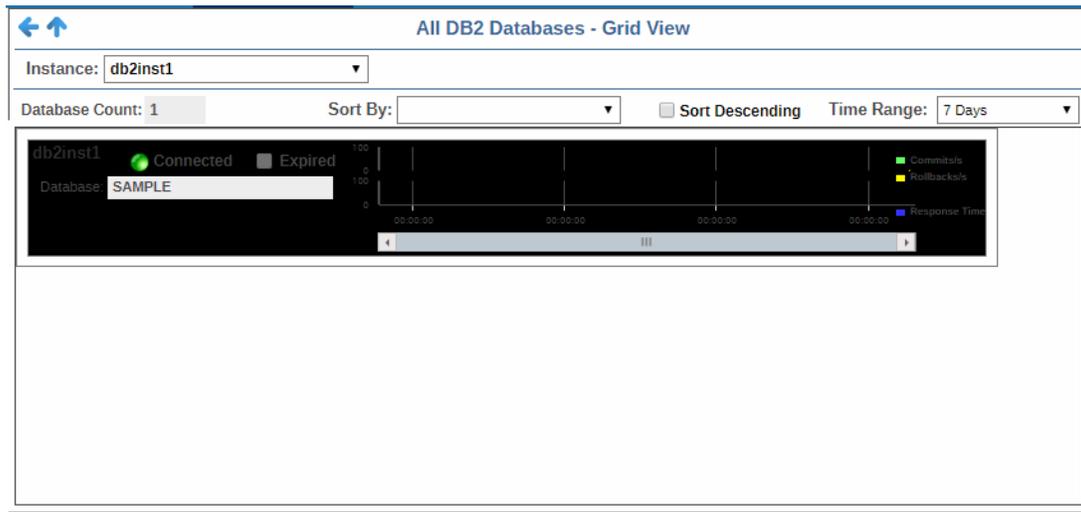
This display provides a grid view of a subset of data in the [“All Databases Table”](#) display.

Each grid object is a different database on the selected instance. Each has a trend graph that traces the **Commits per second**, **Rollbacks per second** and **Response Time** for the database. Choose a **Time Range** for the trend graph to trace, or select **All Data** to include all available data in the trace. Scroll forward and backward in the trend graph and click **Reset** to return to the original state.

Show **Expired** databases first using **Sort By: Expired**. Or sort by instance name, database name or connected state.

Toggle **Sort Descending** to order the grid objects.

Investigate performance metrics details for a database in the “Database Summary” display by clicking on the (left or right side of the) grid object.



**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

- Instance:** Select an instance.
- Database Count:** The number of databases in the display.
- Sort By:** Orders the grid objects as follows:
- **Instance:** Instance names.
  - **Database:** Database names.
  - **Connected:** Databases that are connected.
  - **Expired:** Databases that are in an expired state.
- Sort Descending** Toggle on to order the grid objects.
- Time Range** Choose a time range for the trend graph to trace, from 2 minutes to 7 days, or choose **All Data** to include all available data. The selected time range applies to all grid objects in the display.
- Grid Objects** Each grid object is a different database and values describe each database.
- Connected** The database connection status:
- Red indicates that the database is disconnected.
  - Green indicates that the database is connected.
- Expired** When checked, performance data has not been received within the time specified by your administrator for the **Expire Time**. If your administrator has also set the **Delete Time**, this grid object will be deleted if no data is received within the time specified for deletion.

- Database** The name of the database.
- Trend Graph**
  - **Commits/s** Traces the number of commits per second.
  - **Rollbacks/s** Traces the number of rollbacks per second.
  - **Response Time** Traces the amount of time to respond.

## Single Database

Displays in this View are:

- "Database Summary":
- "Partition Table":
- "Partition Heatmap"
- "Partition Grid"
- "Reports"
- "Trends"
- "DB2 Log"

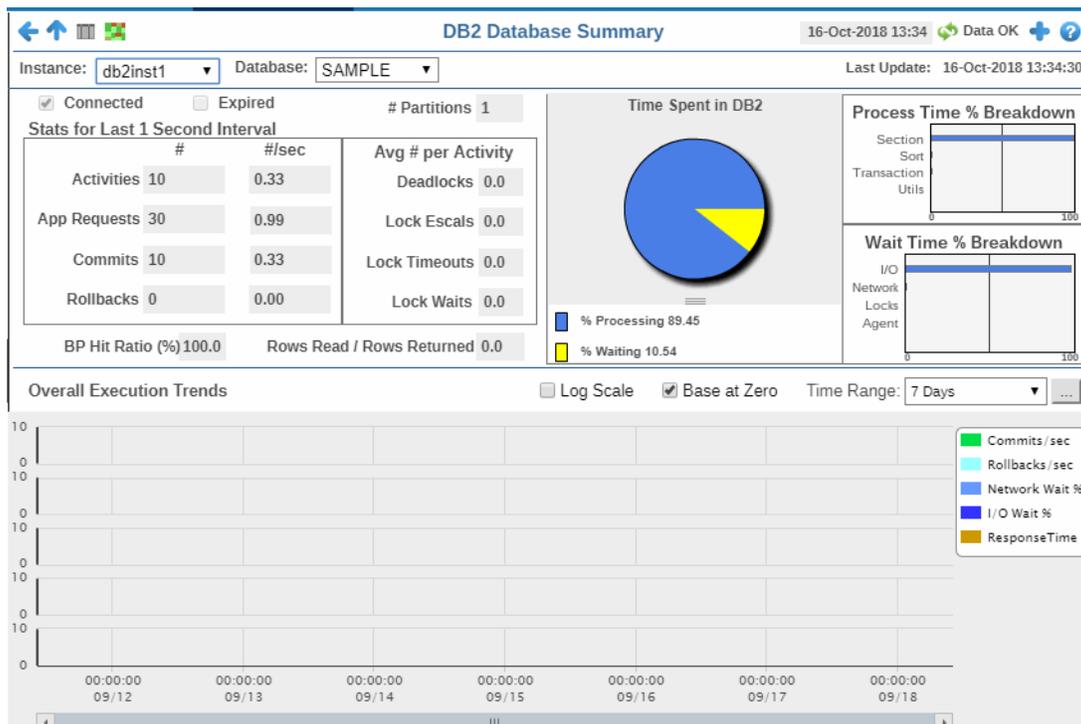
## Database Summary

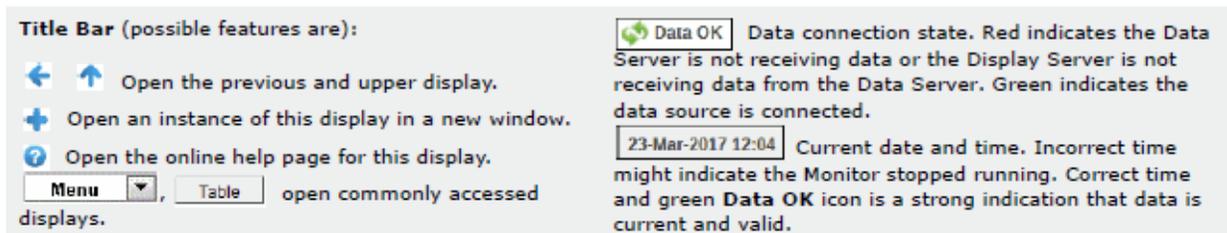
Use this display to investigate the performance and health of a database.

Select an **Instance** and a **Database**. Check main performance statistics such as database load, **Commits per second**, **Processing Time Breakdown** and **Wait Time Breakdown**

View trend graphs tracing **Commits** and **Rollbacks** per second, **Wait** and **Response** times, among others.

Choose a **Time Range** or click  to specify your own.





**Instance:** Choose an instance.

**Database:** Choose a database.

**Connected** When checked, the database is connected.

**Expired** When checked, performance data has not been received within the time specified by your administrator for the **Expire Time**. If your administrator has also set the **Delete Time**, this row will be deleted if no data is received within the time specified for deletion.

**# Partitions** The number of partitions on the database.

#### Stats for Last 1 Second Interval

	#	#/sec
<b>Activites</b>	Total number of activities in the last 1 second.	Value per second.
<b>App Requests</b>	Total number of application requests in the last 1 second.	Value per second.
<b>Commits</b>	Total number of commits in the last 1 second.	Value per second.
<b>Rollbacks</b>	Total number of rollbacks in the last 1 second.	Value per second.

#### Avg # per Activity

<b>Deadlocks</b>	The average number of deadlocks per activity in the last 1 second.
<b>Lock Escal</b>	The average number of lock escalations per activity in the last 1 second.
<b>Lock Timeouts</b>	The average number of lock timeouts per activity in the last 1 second.
<b>Lock Waits</b>	The average number of lock waits per activity in the last 1 second.

#### Time Spent in DB2

**% Processing**

**% Waiting**

- Process Time % Breakdown** Shows the percentage of total processing time for the instance used by the following types of actions:
- Section
  - Sort
  - Transaction
  - Utilities
- Wait Time % Breakdown** Shows the percentage of total wait time used for the instance categorized by:
- Activity
  - IO
  - Network
  - Locks
- BP Hit Ratio %** The current buffer pool hit ratio, which is the total number of pool hits divided by the total number of buffer pool lookups.
- Rows Read / Rows Returned** The number of rows read per number of rows returned.

### Overall ExecutionTrends

For the selected database, trend graphs trace as follows:

- **Commits/sec** Traces the number of commits per second.
- **Rollbacks/sec** Traces the number of rollbacks per second.
- **Network Wait %** Traces the percentage of wait time used by network operations.
- **I/O Wait %** Traces the percentage of wait time used by I/O operations.
- **Response Time** Traces the total amount of wait time.

### Log Scale

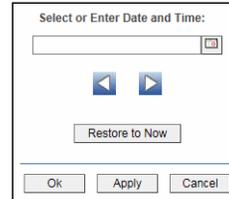
Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.

**Base at Zero**

Select to use zero (**0**) as the Y axis minimum for all graph traces.

**Time Range**

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. By default, the time range end point is the current time. To enter a specific time range, click the associated ellipsis button .



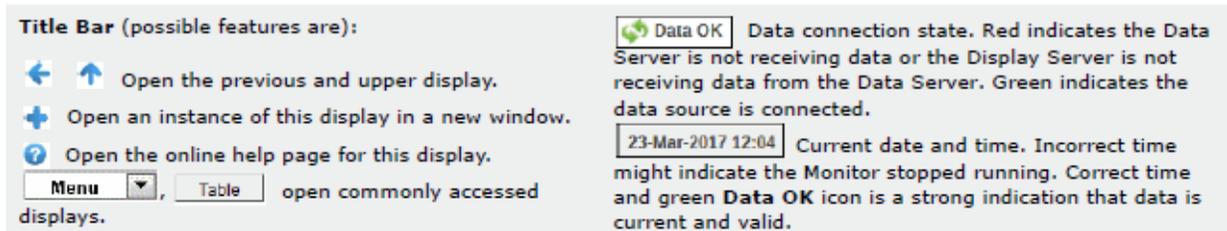
To change the time range click the Open Calendar button , choose the date and time, then click **OK**. Or, enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss** (for example, Aug 21, 2011 12:24 PM) and click **Apply**. Use the Navigation Arrows   to move forward or backward one time period (the time period selected from the Time Range drop-down menu). Click **Restore to Now** to reset the time range end point to the current time.

**Partition Table**

View a list of all partitions on a database, performance metrics for each partition, as well as setup details such as the product associated with the partition and service level.

Use this display to investigate partitioning issues on a database. Select an **Instance** and a **Database**. Sort the list by column values such as **DB2 Status**, **Rollbacks** per second and many others.

DB2 Database Partition Table													
Instance: db2inst1		Database: SAMPLE										18-Oct-2018 10:06 Data OK	
Database Partitions													
Partition #	Expir.	DB2 Status	DB Status	DB Start Time	Local Cons	Remote Cons	App C... Cons	SQL Fa... / sec	Rows Read / sec	Rows Selected / sec	Rows Chan. / sec		
0	<input type="checkbox"/>	ACTIVE	ACTIVE	2018/05/30 10:16:29 PDT	0	11	4	0.00	3.46	4.08	0.0		



### Database Partitions Table

Each row is a different partition on the selected database. Column values describe the partition.

**Instance:** Select an instance.

**Database:** Select a database.

<b>Partition #</b>	The partition number.
<b>Expired</b>	When checked, performance data has not been received within the time specified by your administrator for the <b>Expire Time</b> . If your administrator has also set the <b>Delete Time</b> , this row will be deleted if no data is received within the time specified for deletion.
<b>DB2 Status</b>	The current DB2 status. For example, ACTIVE.
<b>DB Status</b>	The current database status. For example, ACTIVE.
<b>DB Start Time</b>	The date and time the database was last started.
<b>Local Cons</b>	The number of local connections.
<b>Remote Cons</b>	The number of remote connections.
<b>App Cur Cons</b>	The number of currently connected applications.
<b>SQL Faults/sec</b>	The number of SQL faults per second.
<b>Rows Read /sec</b>	The number of rows read per second.
<b>Rows Selected /sec</b>	The number of rows selected per second.
<b>Rows Changed /sec</b>	The number of rows changed per second.
<b>SQL Selects /sec</b>	The number of SQL selects per second.
<b>Commits /sec</b>	The number of commits per second.
<b>Rollbacks /sec</b>	The number of rollbacks per second.
<b>Update/Del/Ins/Stmts / sec</b>	The number of updates, deletions, insertions and statements per second.
<b>Avg Sort Time/Transaction</b>	The average amount of time for sorting transactions.
<b>Product Name</b>	The name of the product.
<b>Service Level</b>	The service level for the product.
<b>Pool Data Hit Ratio</b>	The current buffer pool hit ratio, which is the total number of pool hits divided by the total number of buffer pool lookups.

<b>Pool TempData Hit Ratio</b>	Refer to vendor documentation for details.
<b>Pool TempIndex Hit Ratio</b>	Refer to vendor documentation for details.
<b>Pkg Cache Inserts/K-Trans</b>	Refer to vendor documentation for details.
<b>Lock Wait Time/K-Trans</b>	Refer to vendor documentation for details.
<b>Dirty Steal Triggers/K-Trans</b>	Refer to vendor documentation for details.
<b>Deadlocks &amp; Lock Timeouts/K-Tran</b>	Refer to vendor documentation for details.
<b>Avg Log Write Time / Trans</b>	Refer to vendor documentation for details.
<b>% Agent Usage</b>	The percentage used by agents.
<b>Agents Registered (Top)</b>	The number of registered agents.
<b>Rows Read Returned Ratio</b>	The number of rows read per number of rows returned.
<b>Select %</b>	Refer to vendor documentation for details.
<b>Phys Buffer Pool Read Ratio</b>	Refer to vendor documentation for details.
<b>Phys Buffer Pool Write Ratio</b>	Refer to vendor documentation for details.

## Partition Heatmap

View current alert status and performance metrics of all partitions on a DB2 database.

Answer questions such as, Are any partitions on this database reaching a state of critical health? Do I need to allocate more tablespace to any partitions? Is processing load and number of connections evenly distributed across partitions?

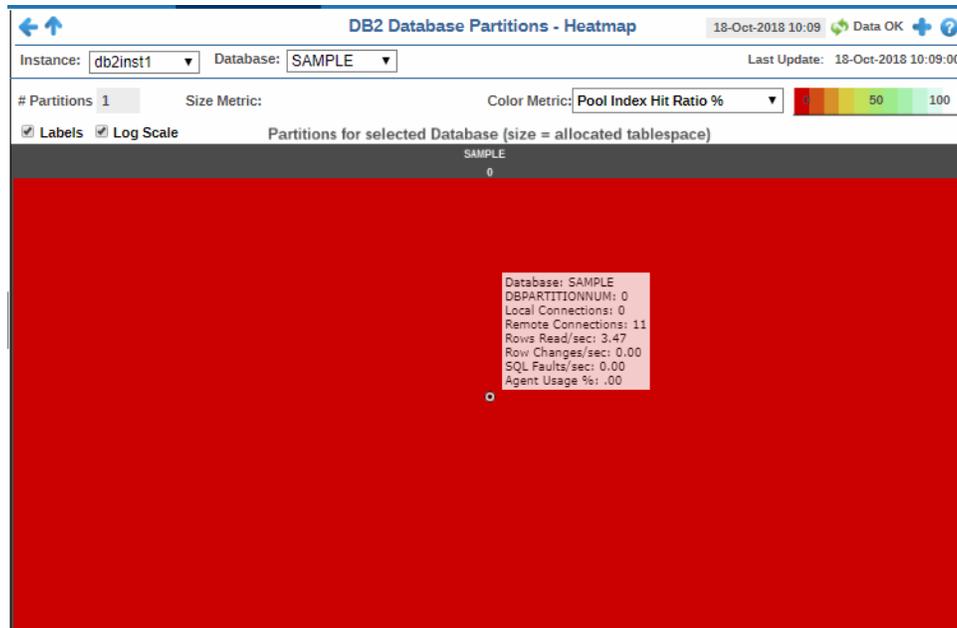
Each rectangle in the heatmap represents a different partition, where the rectangle color indicates the most critical alert state for items associated with that partition, and the rectangle size represents the tablespace allocation size for the partition.

Select an **Instance** and a **Database**. Use the **Metric** drop-down menu to view **Rows Read Per Second**, **Rollbacks Per Second** and **Dirty Steel Triggers/K-Trans**, among many others.

Each metric has its own color gradient bar legend that maps values to colors. By default, the **Commits/sec** metric is shown, which is the number of commits per second for the partition. Values range from **0** to the maximum number in the heatmap, as indicated in the color gradient bar:

-  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
-  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
-  Green indicates that no metrics have exceeded their alert thresholds.

Use the **Labels** check-box  to include or exclude labels in the heatmap, use the **Log Scale** check-box  to apply log scale and mouse over a rectangle to see additional metrics. Click a rectangle to see performance metrics for the database in which the partition resides in the “Database Summary” display.



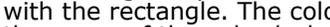
#### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

#### Fields and Data:

- Labels** Select this check box to display the names of the instances at the top of each rectangle in the heatmap.
- Log** Select to this check box to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Auto** Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value.  
**Note:** Some metrics auto-scale automatically, even when **Auto** is not selected.
- Metric** Choose a metric to view in the display. Each rectangle in the heatmap represents a different partition on the selected database. For additional details about the data, refer to vendor documentation.

<b>Commits/sec</b>	<p>The number of application commits per second. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from <b>0</b> to the maximum value in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p> <ul style="list-style-type: none"> <li> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Pool Index Hit Ratio %</b>	<p>The average response time, in milliseconds, for items associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from <b>0</b> to the alert threshold of the <b>Db2ResponseTimeHigh</b> alert. The middle value in the gradient bar indicates the middle value of the range.</p> <ul style="list-style-type: none"> <li> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>SQL Failures/sec</b>	<p>The number of SQL faults per second. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from <b>0</b> to the maximum value in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p> <ul style="list-style-type: none"> <li> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Rows Read /sec</b>	<p>The number of rows read per second. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from <b>0</b> to the maximum value in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p> <ul style="list-style-type: none"> <li> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Rows Selected /sec</b>	<p>The number of rows selected per second. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from <b>0</b> to the maximum value in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p> <ul style="list-style-type: none"> <li> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>

<b>Rows Changed /sec</b>	<p>The number of rows changed per second. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from <b>0</b> to the maximum value in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p> <ul style="list-style-type: none"> <li> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>SQL Select Stmts/ sec</b>	<p>The number of SQL statements selected per second. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from <b>0</b> to the maximum value in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p> <ul style="list-style-type: none"> <li> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Rollbacks/sec</b>	<p>The number of rollbacks per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum number of incoming messages per second. The middle value in the gradient bar indicates the middle value of the range.</p> <ul style="list-style-type: none"> <li> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>SQL Update/Del/ Ins/Stmts /sec</b>	<p>The number of SQL updates, deletions, insertions and statements per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum number of incoming messages per second. The middle value in the gradient bar indicates the middle value of the range.</p> <ul style="list-style-type: none"> <li> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>

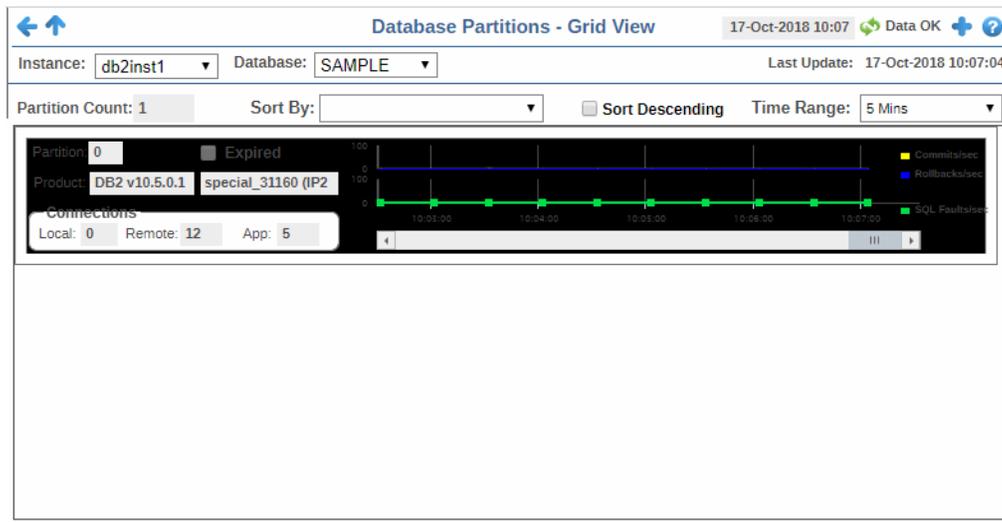
## Partition Grid

This display provides a grid view of a subset of data in the “[Partition Table](#)” display.

Each grid object is a different partition on the selected database. Each has a trend graph that traces the **Commits per second**, **Rollbacks per second** and **SQL Faults per second** for the partition. Choose a **Time Range** for the trend graph to trace, or select **All Data** to include all available data in the trace. Scroll forward and backward in the trend graph and click **Reset** to return to the original state.

Show **Expired** databases first using **Sort By: Expired**. Or sort by partition name. Toggle **Sort Descending** to order the grid objects.

Investigate performance metrics details for a database in the “Database Summary” display by clicking on the (left or right side of the) grid object.



**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

- Instance:** Select an instance.
- Database:** Select a database.
- Partition Count:** The number of partitions in the display.
- Sort By:** Orders the grid objects as follows:
  - **Partition #:** Partition number.
  - **Expired:** Partitions that are in an expired state.
- Sort Descending** Toggle on to order the grid objects.
- Time Range** Choose a time range for the trend graph to trace, from 2 minutes to 7 days, or choose **All Data** to include all available data. The selected time range applies to all grid objects in the display.
- Grid Objects** Each grid object is a different partition and values describe the partition.
  - Partition** The partition number.
  - Expired** When checked, performance data has not been received within the time specified by your administrator for the **Expire Time**. If your administrator has also set the **Delete Time**, this grid object will be deleted if no data is received within the time specified for deletion.
  - Product** The name of the software and version.

- Connections**
  - **Local** The number of current local connections.
  - **Remote** The number of current remote connections.
  - **App** The number of applications
- Trend Graph**
  - **Commits/s** Traces the number of commits per second.
  - **Rollbacks/s** Traces the number of rollbacks per second.
  - **Response Time** Traces the amount of time to respond.

## Reports

Get maximum usage reports for your databases. Select an **Instance**, a **Database**, a **Report** and the number of rows to return (5, 10, 20, 50 or 100 rows). Reports are:

- **Top Applications by Sort Time:** Lists applications with the greatest wait and sort times. Includes User IDs, number of sorts and wait time.
- **Top Lock Waits by Application:** Lists applications with the most lock waits.
- **Top Queries Against Package Cache:** Lists statements with the greatest number of executions against a package cache. Includes wait and total times.
- **Top Queries by Execution Time:** Lists queries with the greatest execution times. Includes elapsed time, AUTH\_ID, AGENT\_ID and SQL\_TEXT.
- **Top Tables by Index Use:** Lists tables containing the most used indexes. Includes Schema, #Leaf Nodes, Nlevels, Index Scans, Key Updates and Page Splits.
- **Top Tables by Rows Selected:** Lists tables containing the most selected rows. Includes Schema, Table Scans, Rows Read, Inserted and Deleted.
- **Top Consuming Transactions:** Lists applications executing transactions that require the greatest CPU use. Includes App Handles, User IDs, wait time, request time, CPU time and client wait time.

Sort reports by any of its columns.

The screenshot shows the 'DB2 Table Access' report interface. At the top, it displays the instance 'db2inst1' and database 'SAMPLE'. The report is titled 'Top Tables by Index Use' and shows a list of tables with their respective index usage statistics. The columns include Schema, Table Name, Index ID, # Leaf Nodes, Nlevels, Index Scans, Key Updates, and Page Splits. The table is sorted by Index Scans in descending order.

Schema	Table Name	Index ID	# Leaf Nodes	Nlevels	Index Scans	Key Updates	Page Splits
SYSTOOLS	HMON_ATM_INFO	1	1	1	640,530	0	0
SYSIBM	SYSTABLES	1	5	2	278,454	0	0
SYSIBM	SYSINDEXES	2	3	2	274,907	0	0
SYSIBM	SYSROUTINES	6	18	2	84,103	0	0
SYSIBM	SYSROUTINES	2	11	2	84,103	0	0
SYSIBM	SYSCOLUMNS	1	82	2	48,045	0	0
SYSIBM	SYSTRIGGERS	2	1	1	44,007	0	0
SYSIBM	SYSODEPROPERTIES	1	1	1	44,007	0	0
SYSTOOLS	POLICY	1	1	1	44,006	0	0
SYSIBM	SYSTASKS	4	1	1	39,926	0	0
SYSIBM	SYSINDEXES	4	2	2	2,073	0	0
SYSIBM	SYSTABLES	6	1	1	1,562	0	0
SYSIBM	SYSPLAN	1	3	2	395	0	0
SYSIBM	SYSPLAN	4	1	1	354	0	0
SYSIBM	SYSDATAPARTITIONS	5	1	1	227	0	0
SYSIBM	SYSCOLDIST	1	279	3	227	0	0
SYSIBM	SYSTABLES	9	1	1	47	0	0
SYSIBM	SYSPLAN	3	2	2	23	0	0
SYSIBM	SYSDBAUTH	1	1	1	13	0	0
SYSIBM	SYSHISTOGRAMTEMPLAT	2	1	1	10	0	0
SYSIBM	SYSWORKLOADCONNATT	1	1	1	4	0	0
SYSIBM	SYSWORKLOADAUTH	1	1	1	4	0	0
SYSIBM	SYSTHRESHOLDS	3	1	1	4	0	0
SYSIBM	SYSSTOGROUPS	2	1	1	4	0	0

**Title Bar (possible features are):**

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu Table open commonly accessed displays.

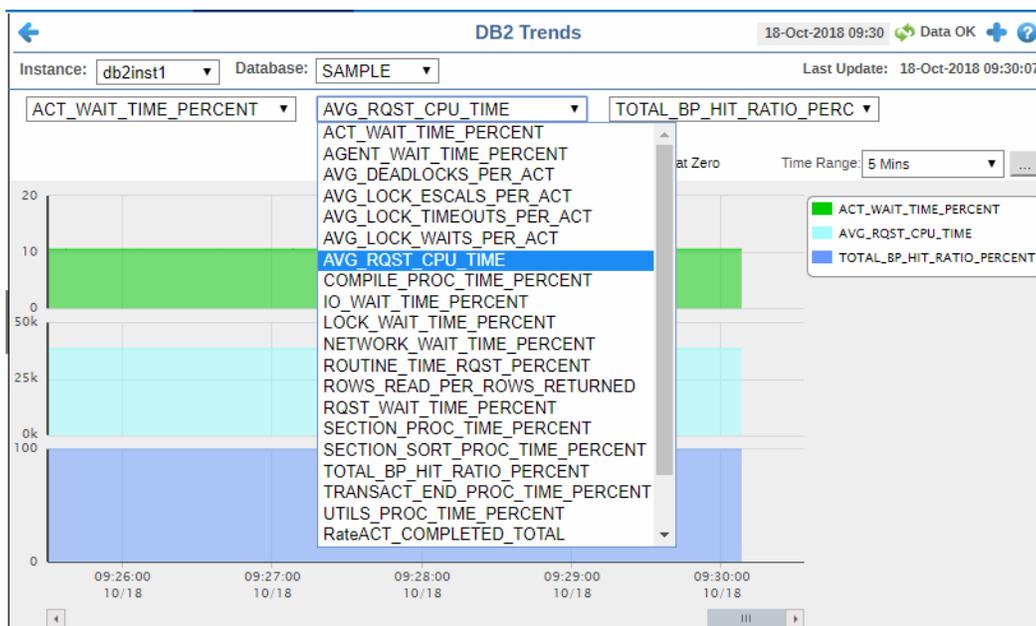
 **Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

 **23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

## Trends

View IBM DB2 database performance and utilization metrics in a trend graph. Select an **Instance**, a **Database** and up to three IBM DB2 metrics. Metric options include (among others) **ACT\_WAIT\_CPU\_TIME**, **AVG\_LOCK\_TIMEOUTS**, **ROWS\_READ\_PER\_ROWS\_RETURNED** and **UTILS\_PROC\_TIME\_PERCENT**.

Choose a **Time Range** or click [...] to specify your own. Mouse-over to see additional details.



**Title Bar (possible features are):**

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu Table open commonly accessed displays.

 **Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

 **23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Log Scale**

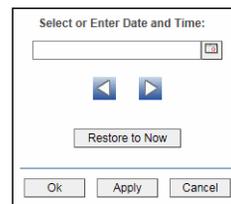
Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.

**Base at Zero**

Select to use zero (0) as the Y axis minimum for all graph traces.

**Time Range**

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. By default, the time range end point is the current time. To enter a specific time range, click the associated ellipsis button .

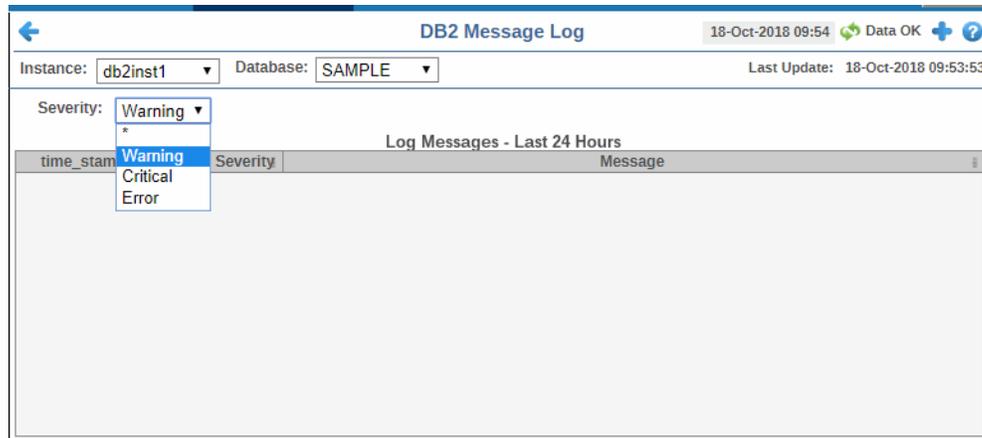


To change the time range click the Open Calendar button , choose the date and time, then click **OK**. Or, enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss** (for example, Aug 21, 2011 12:24 PM) and click **Apply**. Use the Navigation Arrows   to move forward or backward one time period (the time period selected from the Time Range drop-down menu). Click **Restore to Now** to reset the time range end point to the current time.

## DB2 Log

Read database log messages generated in the previous twenty-four hours. Select an **Instance** and a **Database**. Filter the list of log messages using the **Severity** drop-down menu:

- **Warning**: Lists only log messages with a warning severity level.
- **Critical**: Lists only log messages with an alarm severity level.
- **Error**: Lists only log message errors.
- **\*** (asterisk): Lists all log messages.



### Title Bar (possible features are):

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu ↓, Table open commonly accessed displays.

 Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

## IBM DB2 - HTML

The IBM DB2 HTML displays provide extensive visibility into the health and performance of IBM DB2 databases. The HTML version features an ["IBM DB2 Overview Display - HTML"](#) (pictured below), and the following Views which can be found under **Components** tab > **Databases** > **IBM DB2 Database**:

- ["DB2 Instances View - HTML"](#): These displays present metrics about IBM DB2 instances.
- ["DB2 Members View - HTML"](#): These displays present metrics about IBM DB2 members.
- ["DB2 Databases View - HTML"](#): These displays present metrics about IBM DB2 databases.

## IBM DB2 Overview Display - HTML

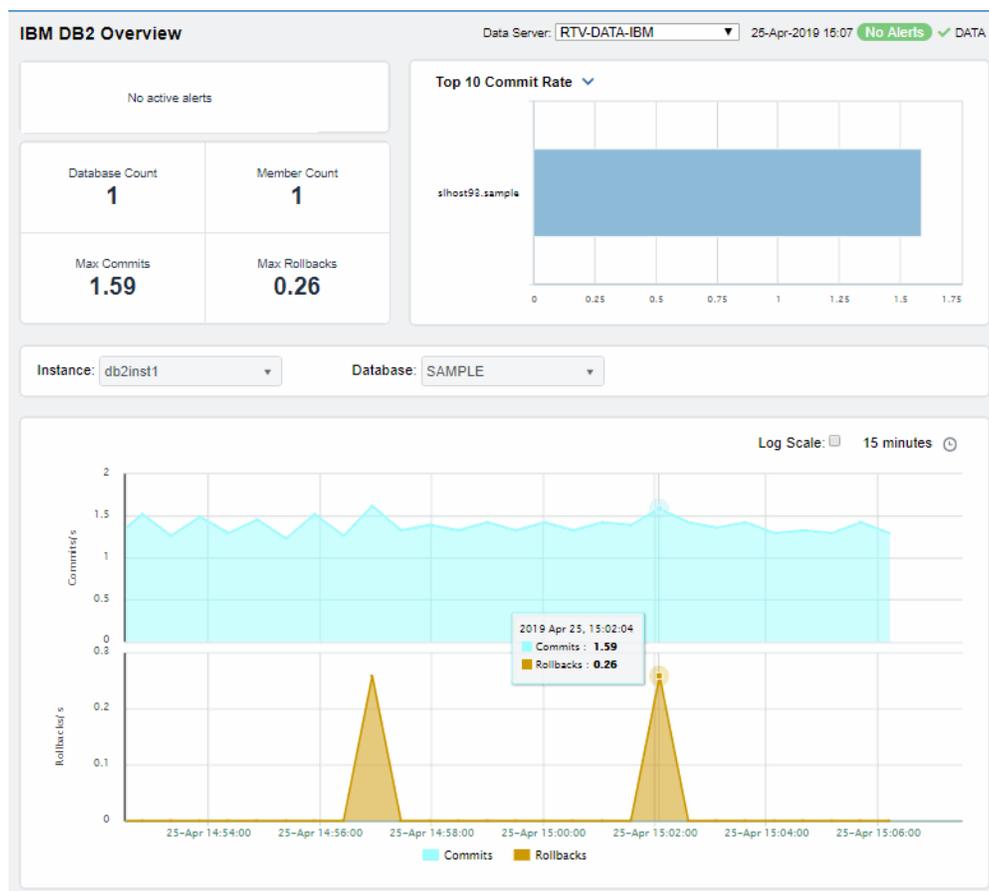
The **IBM DB2 Database Overview** is the top-level display for the IBM DB2 Database Solution Package, which provides a good starting point for immediately getting the status of all your IBM DB2 databases on your Data Server.

You can select the RTView DataServer for which you want to see data and easily view the current data for that DataServer including:

- The total number of **Active Alerts**, including the total number of critical and warning alerts.
- The **Database Count** and **Member Count** for the data server.
- The number of **Max Commits** and **Max Rollbacks** across all databases on the data server.
- A bar graphs showing the **Top 10 Commit Rate**.

You can hover over each region in the upper half of the Overview to see more detail. You can also drill-down to see even more detail in the "[DB2 Members Table - HTML](#)", for example, by clicking on each respective region in the Overview.

The bottom half of the display allows you to select an **Instance** and a **Database** for the performance trend graph to trace **Commits** and **Rollbacks**. You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.



## DB2 Instances View - HTML

Displays in this View are:

- ["DB2 Instances Table - HTML"](#)
- ["DB2 Instances Heatmap - HTML"](#)
- ["Instance Summary - HTML"](#)

### DB2 Instances Table - HTML

Investigate detailed configuration information for all DB2 instances. This display provides a list of all DB2 instances with details about **Partition Counts and Pointer Bit Sizes**, as well as **Fix Packs, Service Levels**, number of **MonitoredDBs** and the most critical **Alert Levels**.

Each row in the table is a different DB2 instance. Click a column header to sort column data in numerical or alphabetical order. Investigate by double-clicking a row to view details for an instance in the ["Instance Summary - HTML"](#) display.

Click right corner of column headers to filter and sort data as well as to choose columns to include in the display.

Instance Name	Alert Level	Alert Count	Partitionable	Release	Fix Pack	Pointer Bit Size	Temp Fix ID	Partition Count	Monitored DBs
db2inst1	OK	0	0	0602010E	1	64	IP23526_31160	1	1

### DB2 Instances Heatmap - HTML

View current alert status and performance metrics of all or just one of your IBM DB2 instances.

Answer questions such as, Are any instances reaching a state of critical health? Is the monitoring load evenly distributed across instances and partitions?

Each rectangle in the heatmap represents a different instance, where the rectangle color indicates the most critical alert state for items associated with that instance and the metric selected. The rectangle size represents the tablespace allocation size for the instance.

By default, the **Alert Severity** metric is shown. Values range from **0 - 2**, as indicated in the color gradient  bar:

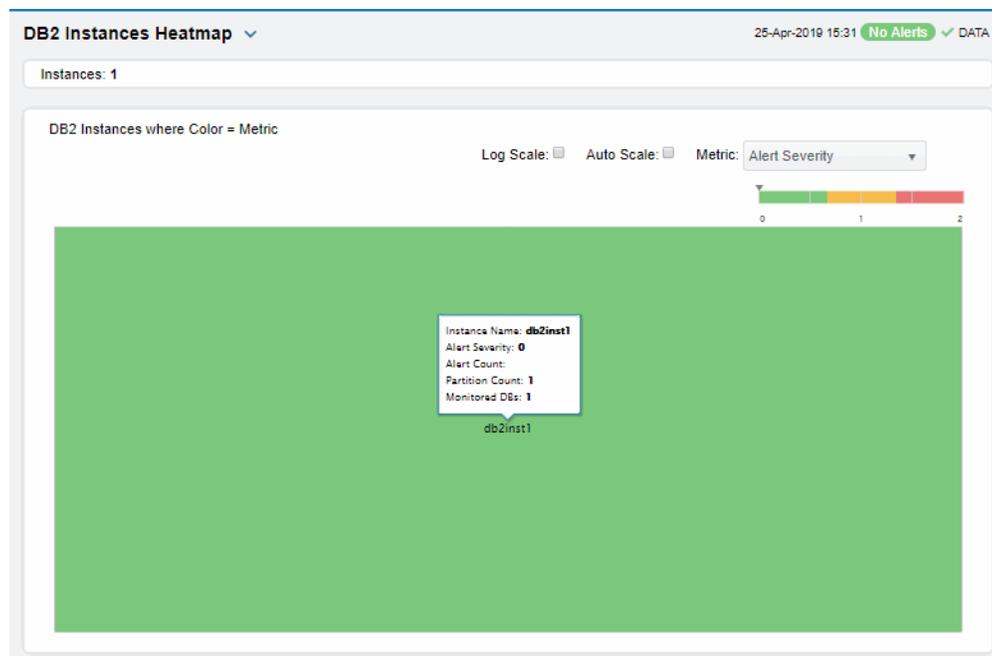
- (2) Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
- (1) Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
- (0) Green indicates that no metrics have exceeded their alert thresholds.

Use the **Metric** drop-down menu to view **Alert Severity**, **Alert Count**, **Partition Count** and **Monitored DBs**.

Mouse-over rectangles to view more details about host performance and status. Toggle between the commonly accessed Table and Heatmap displays by clicking the drop down list on the display title.

You can view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.

Click a rectangle to investigate an instance in the ["Instance Summary - HTML"](#) display.



## Instance Summary - HTML

Select an instance **Instance** to see the following information for that instance:

The **Instance Members** table shows resource allocation and configuration details (CPU Load, Memory) for members (hosts) on a single IBM DB2 instance. Click a row to drill-down to details in the "[DB2 Member Summary - HTML](#)" display.

The **Instance Databases** table shows utilization and processing metrics for databases on a single IBM DB2 instance. Click a row to drill-down to details in the "[Single DB2 Database Summary - HTML](#)" display.

The screenshot displays the 'DB2 Instance Summary' dashboard. At the top, it shows the instance name 'db2inst1' and the date '25-Apr-2019 15:34' with a 'No Alerts' status. Below this, several key metrics are presented in a grid: Release Number (0602010E), Fix Pack (1), Pointer Bits (64), Temp Fix ID (IP23526\_31160), and Service Level (DB2 v10.5.0.1). A Build Level of 'special\_31160' is also shown. The 'Members' section contains a table with columns for Host Name, Expired, Member, Operating Sy., Memory Total, Virtual Memory, CPUs, CPU Load Long, CPU Load Medium, and CPU Load Short. A single row is visible for 'SLHOST93'. The 'Databases' section contains a table with columns for Database Name, Expired, Activities/s, App Requests/s, Commits/s, Rollbacks/s, Agent Wait Time %, IO Wait Time %, Network Wait Time %, and Lock Time. A single row is visible for 'SAMPLE'. At the bottom, it indicates 'Partitionable: 0' and 'Partition Count: 1'.

### Instance Members Table

Each row is a different host member. Column values describe the host.

**Member Count:** The number of members in the table.

**Host Name** The name of the host.

**Expired** When checked, performance data has not been received within the time specified by your administrator for the **Expire Time**. If your administrator has also set the **Delete Time**, this row will be deleted if no data is received within the time specified for deletion.

**Member** The member number.

**OS** The installed operating system.

**CPUs** The number of CPUs and the number of CPUs in use.

<b>CPU Speed</b>	The processor speed.
<b>Total RAM</b>	The total amount of RAM, in megabytes.
<b>Total Virtual Memory</b>	The total amount of virtual memory, in megabytes.
<b>CPU Load Short</b>	Amount of processor load over the short term (defined by the IBM DB2 system, for example, 1-5 minutes).
<b>CPU Load Medium</b>	Amount of processor load over the medium term (defined by the IBM DB2 system, for example, 5-10 minutes).
<b>CPU Load Long</b>	Percentage of CPU load over the long term (defined by the IBM DB2 system, for example, 10-15 minutes).
<b>% CPU Usage</b>	The percentage of CPU used.

**Instance Databases Table**

Each row is a different database. Column values describe the database.

**Database Count:** The number of databases in the table.

<b>Database</b>	The name of the database.
<b>Expired</b>	When checked, performance data has not been received within the time specified by your administrator for the <b>Expire Time</b> . If your administrator has also set the <b>Delete Time</b> , this row will be deleted if no data is received within the time specified for deletion.
<b>Connected</b>	When checked, the database is connected.
<b>Severity</b>	The alert status: <ul style="list-style-type: none"> <li> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Response Time</b>	The response time, in milliseconds.
<b>I/O Wait Time %</b>	The percentage wait time taken by I/O operations.
<b>Network Wait Time %</b>	The percentage wait time taken by the network.
<b>Agent Wait Time %</b>	The percentage wait time taken by agents.
<b>Avg Deadlocks per Activity</b>	The average number of application deadlocks per activity.
<b>Avg Lock Escalations per Activity</b>	The average number of application deadlock escalations per activity.
<b>Avg Lock Timeouts per Activity</b>	The average number of application deadlock timeouts per activity.
<b>Avg Lock Waits per Activity</b>	The average number of application deadlock waits per activity.
<b>Rows Read per Rows Returned</b>	The number of rows read per number of rows returned.
<b>Activites/sec</b>	The number of activities per second.
<b>App Requests/sec</b>	The number of application requests per second.

<b>Commits/sec</b>	The number of application commits per second.
<b>Rollbacks/sec</b>	The number of application rollbacks per second.
<b>Buffer Pool Hit Ratio %</b>	The current buffer pool hit ratio, which is the total number of pool hits divided by the total number of buffer pool lookups.
<b>Activity Wait Time %</b>	The percentage wait time taken by activities.
<b>Avg Request CPU Time</b>	The average amount of CPU time used by requests, in seconds.
<b>Compile Proc Time %</b>	The percentage of time used for compiling processes.
<b>Routine Time Request %</b>	The percentage of time used for routine request processes.
<b>Section Time %</b>	The percentage of time used for section processes.
<b>Section Sort Time %</b>	The percentage of time used for sorting section processes.
<b>BP Hit Ratio %</b>	The current buffer pool hit ratio, which is the total number of pool hits divided by the total number of buffer pool lookups.
<b>Transaction Time %</b>	The percentage of time used for transaction processes.
<b>Utils Proc Time %</b>	The percentage of time used for utilities processes.
<b>Timestamp</b>	The data and time this data was last updated.

## DB2 Members View - HTML

Displays in this View are:

- ["DB2 Members Table - HTML"](#): Get configuration details and utilization metrics for host members of an instance.
- ["DB2 Member Summary - HTML"](#): View trend graph for CPU load and CPU utilization for each member of an instance.

## DB2 Members Table - HTML

Select an instance **Instance** to see a list of all host members of an instance.

Each row is a different host that shows configuration details and utilization metrics for each. Details include OS, RAM, CPU load (short, medium, long) and virtual memory.

Click a column header to sort column data in numerical or alphabetical order. Drill-down and investigate by clicking a row to view details for a host member in the ["DB2 Member Summary - HTML"](#) display.

DB2 Members Table 25-Apr-2019 15:36 ✓ DATA

Instance:  Members: 1

Host Name	Expired	Member	Operating Sy.	Memory Total	Virtual Memory	CPUs	CPU Load Long	CPU Load Medium	CPU Load Short	CPU
SLHOST93		0	Linux 2.6.32(x86_64)	3835	7867	1 of 1 in use	0.28	0.30	0.31	

### Instance Members Table

Each row is a different host member. Column values describe the host.

**Member Count:** The number of members in the table.

**Host Name** The name of the host.

**Expired** When checked, performance data has not been received within the time specified by your administrator for the **Expire Time**. If your administrator has also set the **Delete Time**, this row will be deleted if no data is received within the time specified for deletion.

**Instance** The name of the instance that the host is a member of.

**Member** The name of the member.

**OS** The installed operating system.

**CPUs** The number of CPUs and the number of CPUs in use.

**CPU Speed** The processor speed.

**Total RAM** The total amount of RAM, in megabytes.

**Total Virtual Memory** The total amount of virtual memory, in megabytes.

**CPU Load Short** Amount of processor load over the short term (defined by the IBM DB2 system, for example, 1-5 minutes).

**CPU Load Medium** Amount of processor load over the medium term (defined by the IBM DB2 system, for example, 5-10 minutes).

- CPU Load Long**                      Percentage of CPU load over the long term (defined by the IBM DB2 system, for example, 10-15 minutes).
- % CPU Usage**                        The percentage of CPU used.
- Timestamp**                            The data and time this data was last updated.

## DB2 Member Summary - HTML

Use this display to study trends for a single host member, such as CPU usage and loads (Short, Medium and Long).

Select an **Instance** and a **Member** from the drop-down menus. Check status of the selected host member **TableSpace** allocations and current **Buffer Pool Hit Ratios**.

Choose a **Time Range** or click  to specify your own. Mouse-over to see additional details..



**Instance:** Choose an instance.

**Member:** Choose a member.

**Expired** When checked, performance data has not been received within the time specified by your administrator for the **Expire Time**.  
If your administrator has also set the **Delete Time**, this row will be deleted if no data is received within the time specified for deletion.

**Host Name** The name of the selected host member.

**OS** The operating system on the selected host member.

#### **Buffer Pool Hit Ratios**

Each row in the table in a different buffer pool. Values describe values for the selected host member.

**BP\_NAME** The name of the buffer pool.

**Data**

**Index**

**Xda**

#### **TableSpace**

Values describe TableSpace values for the selected host member.

**TableSpace** The name of the TableSpace.

**% Utilization** The percentage of TableSpace used.

**Allocated** The amount of TableSpace allocated.

**Free(MB)** The amount of free TableSpace, in megabytes.

**Auto Extend** Indicates whether auto extend is enabled. Y/N

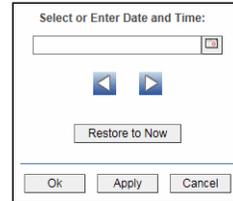
#### **Trend Graph**

For the selected host member, trend graphs trace as follows:

- **CPU Load Short** Traces the amount of CPU load over the short term (defined by the IBM DB2 system, for example, 1-5 minutes).
- **CPU Load Medium** Traces the amount of CPU load over the medium term (defined by the IBM DB2 system, for example, 5-10 minutes).
- **CPU Load Long** Traces the amount of CPU load over the long term (defined by the IBM DB2 system, for example, 10-15 minutes).
- **CPU Usage %** Traces the percentage of CPU used.

**Log Scale** Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.

- Base at Zero** Select to use zero (**0**) as the Y axis minimum for all graph traces.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. By default, the time range end point is the current time. To enter a specific time range, click the associated ellipsis button .



To change the time range click the Open Calendar button , choose the date and time, then click **OK**. Or, enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss** (for example, Aug 21, 2011 12:24 PM) and click **Apply**. Use the Navigation Arrows   to move forward or backward one time period (the time period selected from the Time Range drop-down menu). Click **Restore to Now** to reset the time range end point to the current time.

## DB2 Databases View - HTML

Displays in this View are:

- ["DB2 Databases Table - HTML"](#)
- ["Single DB2 Database Summary - HTML"](#)
- ["Partition Table"](#)
- ["Partition Heatmap"](#)

### DB2 Databases Table - HTML

Select an **Instance** and view a list of all its databases and their performance metrics.

Use this display to identify which databases are having or causing issues for the instance.

Sort the list of databases by **Alert Severity** or by any other table column value.

In the bar graphs, view **Wait Time Breakdown** for the instance (by Activity, IO, Network and Lock), as well as **Processing Time Breakdown** (by Compile, Section, Sort, Transaction and Utilities) for the instance.

Click a row to drill-down to details in the “Single DB2 Database Summary - HTML” display.

DB2 Databases Table 25-Apr-2019 15:42 ✓ DATA

Instance:  Databases: 1

Database Name	Expired	Activities/s	App Requests/s	Commits/s	Rollbacks/s	Agent Wait Time %	IO Wait Time %	Network Wait Time %	Lock Wait Time %
db2inst1.SAMPLE		1.20	3.69	1.20	0.00	0.00	98.48	0.38	0.08

### Instance Databases Table

Each row is a different database on the selected instance. Column values describe the database.

**Database Count:** The number of databases in the table.

**Database**

The name of the database.

**Expired**

When checked, performance data has not been received within the time specified by your administrator for the **Expire Time**. If your administrator has also set the **Delete Time**, this row will be deleted if no data is received within the time specified for deletion.

**Connected**

When checked, the database is connected.

**Severity**

The alert status:

● Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.

● Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.

● Green indicates that no metrics have exceeded their alert thresholds.

**Response Time**

The response time, in milliseconds.

**I/O Wait Time %**

The percentage wait time taken by I/O operations.

**Network Wait Time %**

The percentage wait time taken by the network.

**Agent Wait Time %**

The percentage wait time taken by agents.

**Avg Deadlocks per Activity**

The average number of application deadlocks per activity.

<b>Avg Lock Escalations per Activity</b>	The average number of application deadlock escalations per activity.
<b>Avg Lock Timeouts per Activity</b>	The average number of application deadlock timeouts per activity.
<b>Avg Lock Waits per Activity</b>	The average number of application deadlock waits per activity.
<b>Rows Read per Rows Returned</b>	The number of rows read per number of rows returned.
<b>Activites/sec</b>	The number of activities per second.
<b>App Requests/sec</b>	The number of application requests per second.
<b>Commits/sec</b>	The number of application commits per second.
<b>Rollbacks/sec</b>	The number of application rollbacks per second.
<b>Buffer Pool Hit Ratio %</b>	The current buffer pool hit ratio, which is the total number of pool hits divided by the total number of buffer pool lookups.
<b>Activity Wait Time %</b>	The percentage wait time taken by activities.
<b>Avg Request CPU Time</b>	The average amount of CPU time used by requests, in seconds.
<b>Compile Proc Time %</b>	The percentage of time used for compiling processes.
<b>Routine Time Request %</b>	The percentage of time used for routine request processes.
<b>Section Time %</b>	The percentage of time used for section processes.
<b>Section Sort Time %</b>	The percentage of time used for sorting section processes.
<b>BP Hit Ratio %</b>	The current buffer pool hit ratio, which is the total number of pool hits divided by the total number of buffer pool lookups.
<b>Transaction Time %</b>	The percentage of time used for transaction processes.
<b>Utils Proc Time %</b>	The percentage of time used for utilities processes.
<b>Activity Time (ms)</b>	The amount of time, in milliseconds, used for activity processes.
<b>Agent Idle Time (ms)</b>	The amount of time, in milliseconds, that agents were in an idle state.
<b>CPU Time (ms)</b>	The amount of time, in milliseconds, used for CPU processes.
<b>Processing Time (ms)</b>	The amount of time, in milliseconds, used for processing requests.
<b>Request Time (ms)</b>	The amount of time, in milliseconds, used for processing requests.
<b>DB Processing %</b>	The percentage of time used for database processes.
<b>DB Wait %</b>	The percentage of time used for database waits.
<b># Connections</b>	The current number of database connections.
<b>Timestamp</b>	The data and time this data was last updated.

**Wait Time Breakdown**

Shows the percentage of total wait time used for the instance categorized by:

- Activity
- IO
- Network
- Lock

**Processing Time Breakdown**

Shows the percentage of total processing time for the instance used by the following types of actions:

- Compile
- Section
- Sort
- Transaction
- Utilities

**Single DB2 Database Summary - HTML**

Use this display to investigate the performance and health of a database.

Select an **Instance** and a **Database**. Check main performance statistics such as database load, **Commits per second**, **Processing Time Breakdown** and **Wait Time Breakdown**

View trend graphs tracing **Commits** and **Rollbacks** per second, **Wait** and **Response** times, among others.

Choose a **Time Range** or click  to specify your own.



**Instance:** Choose an instance.

**Database:** Choose a database.

**Connected** When checked, the database is connected.

**Expired** When checked, performance data has not been received within the time specified by your administrator for the **Expire Time**. If your administrator has also set the **Delete Time**, this row will be deleted if no data is received within the time specified for deletion.

**# Partitions** The number of partitions on the database.

**Stats for Last 1 Second Interval**

#

#/sec

<b>Activites</b>	Total number of activities in the last 1 second.	Value per second.
<b>App Requests</b>	Total number of application requests in the last 1 second.	Value per second.
<b>Commits</b>	Total number of commits in the last 1 second.	Value per second.
<b>Rollbacks</b>	Total number of rollbacks in the last 1 second.	Value per second.

**Avg # per Activity**

<b>Deadlocks</b>	The average number of deadlocks per activity in the last 1 second.
<b>Lock Escal</b>	The average number of lock escalations per activity in the last 1 second.
<b>Lock Timeouts</b>	The average number of lock timeouts per activity in the last 1 second.
<b>Lock Waits</b>	The average number of lock waits per activity in the last 1 second.

**Time Spent in DB2****% Processing****% Waiting****Process Time % Breakdown**

Shows the percentage of total processing time for the instance used by the following types of actions:

- Section
- Sort
- Transaction
- Utilities

**Wait Time % Breakdown**

Shows the percentage of total wait time used for the instance categorized by:

- Activity
- IO
- Network
- Locks

**BP Hit Ratio %**

The current buffer pool hit ratio, which is the total number of pool hits divided by the total number of buffer pool lookups.

**Rows Read / Rows Returned**

The number of rows read per number of rows returned.

**Overall ExecutionTrends**

For the selected database, trend graphs trace as follows:

- **Commits/sec** Traces the number of commits per second.
- **Rollbacks/sec** Traces the number of rollbacks per second.
- **Network Wait %** Traces the percentage of wait time used by network operations.
- **I/O Wait %** Traces the percentage of wait time used by I/O operations.
- **Response Time** Traces the total amount of wait time.

**Log Scale**

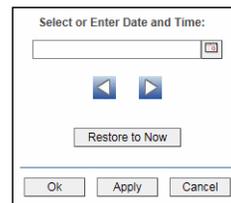
Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.

**Base at Zero**

Select to use zero (0) as the Y axis minimum for all graph traces.

**Time Range**

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. By default, the time range end point is the current time. To enter a specific time range, click the associated ellipsis button .



To change the time range click the Open Calendar button , choose the date and time, then click **OK**. Or, enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss** (for example, Aug 21, 2011 12:24 PM) and click **Apply**. Use the Navigation Arrows   to move forward or backward one time period (the time period selected from the Time Range drop-down menu). Click **Restore to Now** to reset the time range end point to the current time.

**Partition Table - HTML**

View a list of all partitions on a database, performance metrics for each partition, as well as setup details such as the product associated with the partition and service level.

Use this display to investigate partitioning issues on a database. Select an **Instance** and a **Database**. Sort the list by column values such as **DB2 Status**, **Rollbacks** per second and many others.

Partition #	Expired	DB2 Status	DB Status	DB Start Time	Local Cons	Remote Cons	App Cur Cons	SQL Faults/s	Rows Re
0	<input type="checkbox"/>	ACTIVE	ACTIVE	Tue Apr 23 2019 11:01:09 GMT	0	13	6	0.000	1

### Database Partitions Table

Each row is a different partition on the selected database. Column values describe the partition.

**Instance:** Select an instance.

**Database:** Select a database.

<b>Partition #</b>	The partition number.
<b>Expired</b>	When checked, performance data has not been received within the time specified by your administrator for the <b>Expire Time</b> . If your administrator has also set the <b>Delete Time</b> , this row will be deleted if no data is received within the time specified for deletion.
<b>DB2 Status</b>	The current DB2 status. For example, ACTIVE.
<b>DB Status</b>	The current database status. For example, ACTIVE.
<b>DB Start Time</b>	The date and time the database was last started.
<b>Local Cons</b>	The number of local connections.
<b>Remote Cons</b>	The number of remote connections.
<b>App Cur Cons</b>	The number of currently connected applications.
<b>SQL Faults/sec</b>	The number of SQL faults per second.
<b>Rows Read /sec</b>	The number of rows read per second.

<b>Rows Selected /sec</b>	The number of rows selected per second.
<b>Rows Changed /sec</b>	The number of rows changed per second.
<b>SQL Selects /sec</b>	The number of SQL selects per second.
<b>Commits /sec</b>	The number of commits per second.
<b>Rollbacks /sec</b>	The number of rollbacks per second.
<b>Update/Del/Ins/Stmts / sec</b>	The number of updates, deletions, insertions and statements per second.
<b>Avg Sort Time/Transaction</b>	The average amount of time for sorting transactions.
<b>Product Name</b>	The name of the product.
<b>Service Level</b>	The service level for the product.
<b>Pool Data Hit Ratio</b>	The current buffer pool hit ratio, which is the total number of pool hits divided by the total number of buffer pool lookups.
<b>Pool TmpData Hit Ratio</b>	Refer to vendor documentation for details.
<b>Pool TmpIndex Hit Ratio</b>	Refer to vendor documentation for details.
<b>Pkg Cache Inserts/K-Trans</b>	Refer to vendor documentation for details.
<b>Lock Wait Time/K-Trans</b>	Refer to vendor documentation for details.
<b>Dirty Steal Triggers/K-Trans</b>	Refer to vendor documentation for details.
<b>Deadlocks &amp; Lock Timeouts/K-Tran</b>	Refer to vendor documentation for details.
<b>Avg Log Write Time / Trans</b>	Refer to vendor documentation for details.
<b>% Agent Usage</b>	The percentage used by agents.
<b>Agents Registered (Top)</b>	The number of registered agents.
<b>Rows Read Returned Ratio</b>	The number of rows read per number of rows returned.
<b>Select %</b>	Refer to vendor documentation for details.
<b>Phys Buffer Pool Read Ratio</b>	Refer to vendor documentation for details.
<b>Phys Buffer Pool Write Ratio</b>	Refer to vendor documentation for details.

## Partition Heatmap - HTML

View current alert status and performance metrics of all partitions on a DB2 database.

Answer questions such as, Are any partitions on this database reaching a state of critical health? Do I need to allocate more tablespace to any partitions? Is processing load and number of connections evenly distributed across partitions?

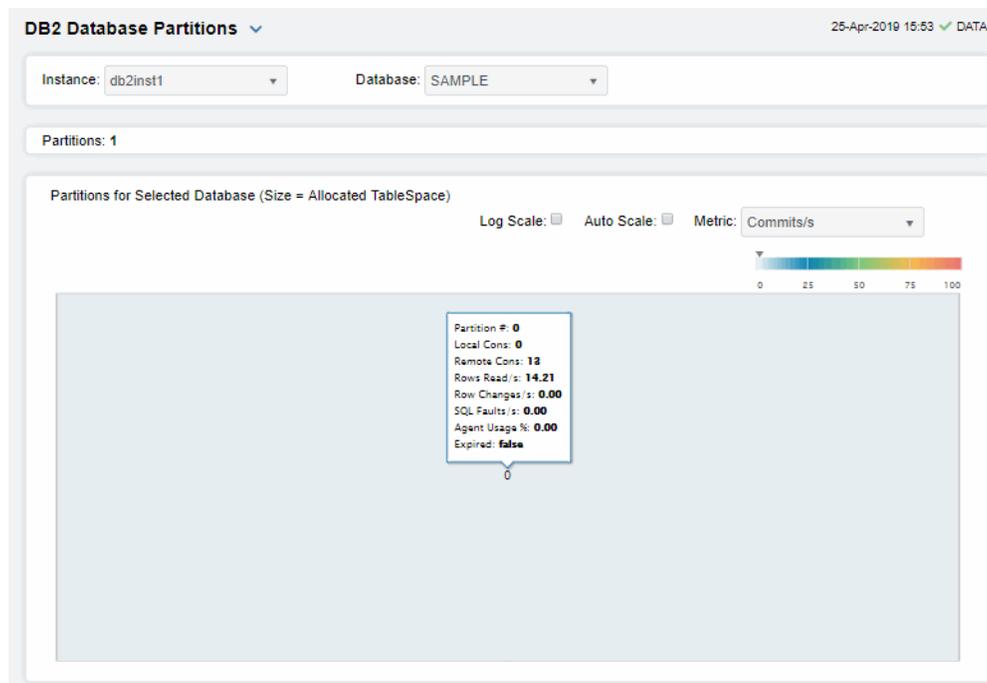
Each rectangle in the heatmap represents a different partition, where the rectangle color indicates the most critical alert state for items associated with that partition, and the rectangle size represents the tablespace allocation size for the partition.

Select an **Instance** and a **Database**. Use the **Metric** drop-down menu to view **Rows Read Per Second**, **Rollbacks Per Second** and **Dirty Steel Triggers/K-Trans**, among many others.

Each metric has its own color gradient bar legend that maps values to colors. By default, the **Commits/sec** metric is shown, which is the number of commits per second for the partition. Values range from **0** to the maximum number in the heatmap, as indicated in the color gradient bar:

- Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
- Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
- Green indicates that no metrics have exceeded their alert thresholds.

Use the **Labels** check-box  to include or exclude labels in the heatmap, use the **Log Scale** check-box  to apply log scale and mouse over a rectangle to see additional metrics. Click a rectangle to see performance metrics for the database in which the partition resides in the ["Single DB2 Database Summary - HTML"](#) display.



### Fields and Data:

- Labels** Select this check box to display the names of the instances at the top of each rectangle in the heatmap.
- Log** Select to this check box to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Auto** Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value.
- Note:** Some metrics auto-scale automatically, even when **Auto** is not selected.

**Metric** Choose a metric to view in the display. Each rectangle in the heatmap represents a different partition on the selected database. For additional details about the data, refer to vendor documentation.

**Commits/sec**

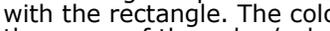
The number of application commits per second. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the maximum value in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

 Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.

 Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.

 Green indicates that no metrics have exceeded their alert thresholds.

**Pool Index Hit Ratio %**

The average response time, in milliseconds, for items associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of the **Db2ResponseTimeHigh** alert. The middle value in the gradient bar indicates the middle value of the range.

 Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.

 Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.

 Green indicates that no metrics have exceeded their alert thresholds.

**SQL Failures/sec**

The number of SQL faults per second. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the maximum value in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

 Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.

 Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.

 Green indicates that no metrics have exceeded their alert thresholds.

**Rows Read /sec**

The number of rows read per second. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the maximum value in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

 Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.

 Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.

 Green indicates that no metrics have exceeded their alert thresholds.

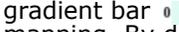
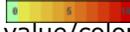
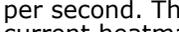
**Rows Selected /sec**

The number of rows selected per second. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the maximum value in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

 Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.

 Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.

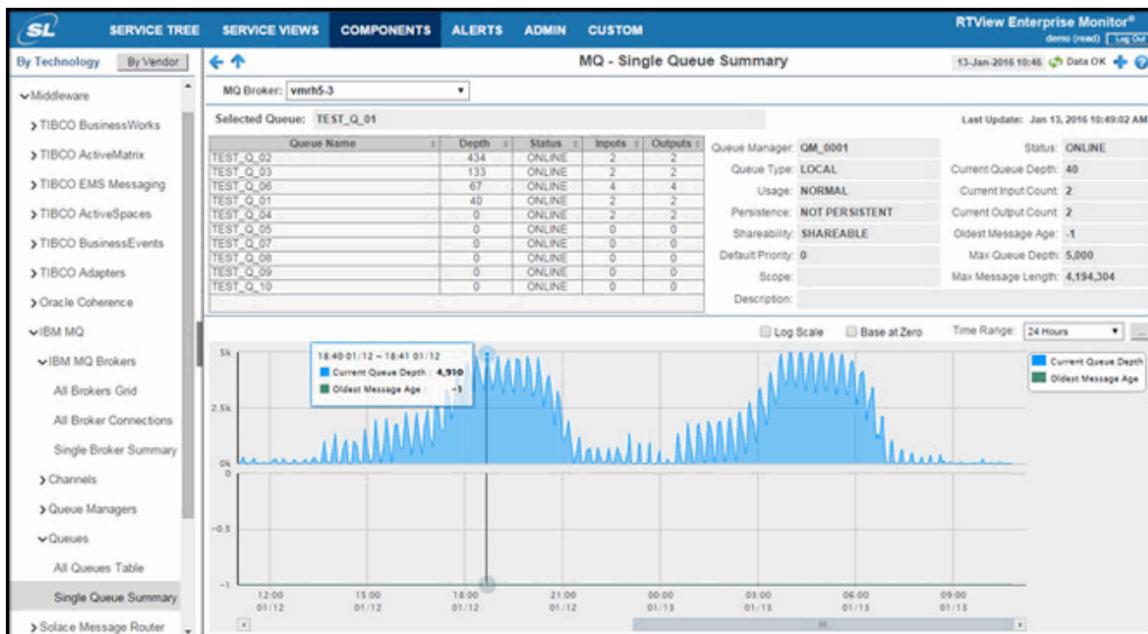
 Green indicates that no metrics have exceeded their alert thresholds.

<b>Rows Changed /sec</b>	<p>The number of rows changed per second. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from <b>0</b> to the maximum value in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p> <ul style="list-style-type: none"> <li> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>SQL Select Stmts/sec</b>	<p>The number of SQL statements selected per second. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from <b>0</b> to the maximum value in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p> <ul style="list-style-type: none"> <li> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Rollbacks/sec</b>	<p>The number of rollbacks per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum number of incoming messages per second. The middle value in the gradient bar indicates the middle value of the range.</p> <ul style="list-style-type: none"> <li> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>SQL Update/Del/Ins/Stmts /sec</b>	<p>The number of SQL updates, deletions, insertions and statements per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum number of incoming messages per second. The middle value in the gradient bar indicates the middle value of the range.</p> <ul style="list-style-type: none"> <li> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>

## IBM MQ

Gain real-time visibility into the health and performance of IBM® MQ objects including brokers, queues, channels and queue managers. The following IBM MQ Monitor Views can be found under **Components** tab > **Middleware**> **IBM MQ**:

- **"IBM MQ Brokers View"**: The displays in this View present performance and utilization metrics for your IBM MQ Brokers.
- **"Channels View"**: The displays in this View present performance and utilization metrics for your IBM MQ Channels.
- **"Queue Managers"**: The displays in this View present performance and utilization metrics for your IBM MQ Queue Managers.
- **"Queues"**: The displays in this View present performance and utilization metrics for your IBM MQ Queues.



## Overview

RTView Enterprise and the Solution Package for IBM® MQ provide out of box performance and availability monitoring for support teams and IBM MQ administrators. Configuration options enable both consolidated views across the enterprise or views configured for specific support teams. As part of an end to end monitoring solution, users can view IBM MQ performance in the context of an application or service. This provides visibility into how IBM MQ performance is impacting adjacent technologies and the resulting business impact. Typical installations of RTView Enterprise and its solution packages take only a few hours, while developing custom views for a variety of IT and development roles can be achieved in just days.

## Key Features

- Monitor real-time performance for early warning
- Analyze historical performance to differentiate trends and spikes
- Out of the box discovery and monitoring of key metrics
- Powerful diagnostics and correlations for complex performance analysis
- View IBM MQ in an application context for Application Support teams and Operations
- Minimal training, highly configurable by business and technical users

## Metrics for IBM MQ

- **All MQ Brokers:**

Queue Manager Status, Number of Channels, Number of Queues, Total Queue Depth

- **MQ Broker Summary:**

Total Queue Depth

Overall Health State: Queue Manager, Channel, Queue Depth High, Queue Full

- **MQ Broker Connections:**

Connection Status, Alert Status, Channel, Model Queue Name, Max Retries, Retry Interval, Wait Interval, Connection

- **All Queues:**

Queue Manager, Queue type, Status, Alert State, Outputs, Inputs, Depth, Max Depth, Persistence Settings, Description, Max Message Length, Host, Default Priority, Get Messages, Put Messages, Scope, Shareability, Usage, Connection

- **Prebuilt Displays:**

All Brokers Grid, Single Broker Summary, All Brokers detail table

All Channels table, Single Channel Summary, Single Channel Detail

All Queue Managers Detail table

All Queues Table, Single Queue Summary

All Trend Graphs show Historical Data

## End-to-End Context for IBM MQ

- Custom flow diagrams help visualize complex applications and IBM MQ's place in that architecture
- Provides an Intuitive View of How IBM MQ Interacts with other Enterprise PaaS Components
- Designed and Developed for Large Scale, Mission Critical Environments

The following WebSphere Monitor Views (and their associated displays) can be found under **Components** tab > **Middleware**> **IBM MQ**.

**Note:** This document assumes familiarity with IBM Websphere MQ. For details about IBM Websphere MQ, refer to vendor documentation.

This section contains the following:

- **"IBM MQ Brokers View"**: The displays in this View present performance and utilization metrics for your IBM MQ Brokers.
- **"Channels View"**: The displays in this View present performance and utilization metrics for your IBM MQ Channels.
- **"Queue Managers"**: The displays in this View present performance and utilization metrics for your IBM MQ Queue Managers.
- **"Queues"**: The displays in this View present performance and utilization metrics for your IBM MQ Queues.

## IBM MQ Brokers View

See performance and utilization metrics for all of your IBM MQ Brokers.

Displays in this View are:

- **"All Brokers Grid"**: This display presents a high-level perspective of utilization metrics for each IBM MQ Broker.
- **"All Broker Connections"**: This display presents detailed connection metrics for each IBM MQ Broker.
- **"Single Broker Summary"**: This display presents performance metrics for a single IBM MQ Broker, as well as detailed metrics for its channels and queues.

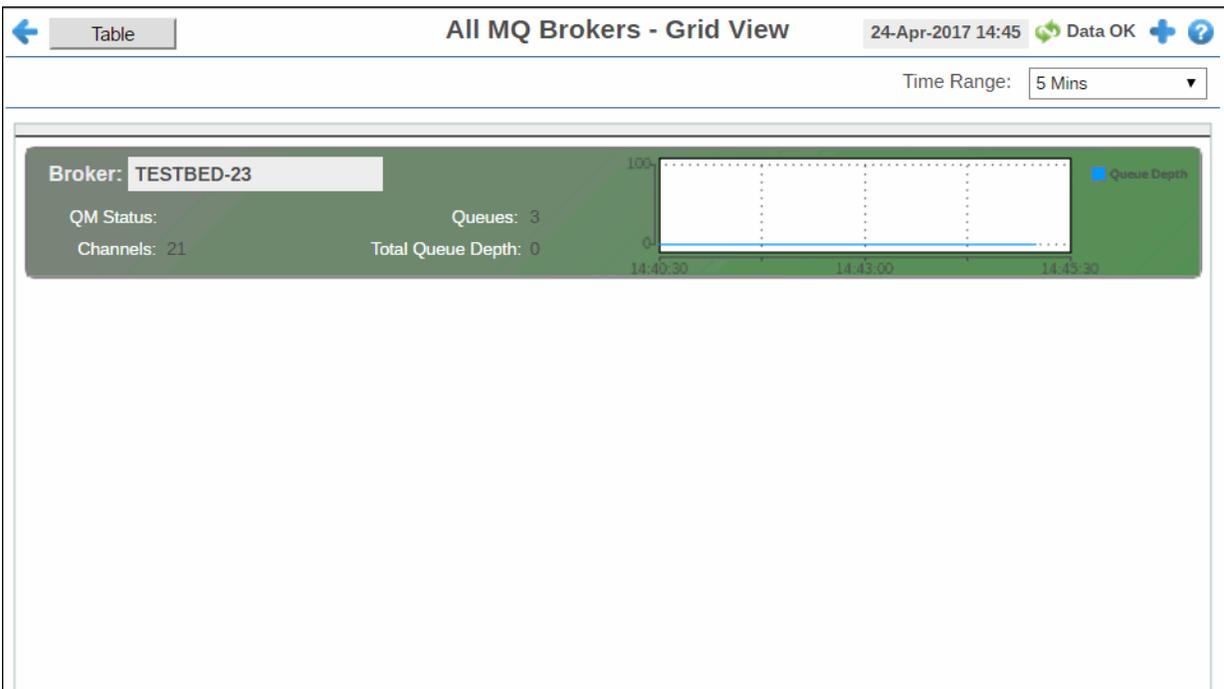
### All Brokers Grid

Track current and historical utilization and performance trends of all MQ Brokers in parallel. Use this display to quickly identify hosts with performance issues and verify whether the load is distributed evenly across brokers.

Each grid object is a different MQ Broker. Inactive brokers are shown in dark red, active brokers are shown in green. Metrics include QM status, queue depth and the number of channels per broker.

This display contains data obtained from IBM MQ. For example, MQIACH\_HB\_INTERVAL, MQIA\_MONITORING\_CHANNEL and MQIACH\_MSG\_COMPRESSION. For details, refer to vendor documentation.

Choose a time range to display from the drop-down menu and drill-down and investigate by clicking a broker to view details in the "Single Broker Summary" display.



**Title Bar** (possible features are):

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu ▾, Table open commonly accessed displays.

 **Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

 **23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**.

**Fields and Data:**

- Broker** The name of the broker.
- QM Status:** The status of the Queue Manager on the broker.
- Queues** The number of queues on the broker.
- Channels** The number of channels on the broker.
- Total Queue Depth** The total queue depth.
- Trend Graph** Traces the queue depth on the broker.

## All Broker Connections

View connection performance details for each MQ Broker in a tabular format, including current status, greatest alert severity and wait time to connect. Each row in the table is a different MQ Broker. Inactive brokers are shown in dark red.

Drill-down and investigate by clicking a row in the table to view details for the selected connection in the “[Single Broker Summary](#)” display.

Broker	Host	Port	Connected	Alert	Channel	Model Queue Name
TESTBED-23	192.168.200.73	1417			ADMIN.TEST.SVRCONN	SYSTEM.DEFAULT.MODE

### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

### Table

Each table row is a different connection. Column values describe the connection except where noted.

- Broker** The name of the broker.
- Host** The name of the host.
- Port** The port number used.
- Connected** The current connection state:
  - Disconnected
  - Connected

<b>Alert</b>	The current alert severity:  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
<b>Channel</b>	The name of the channel.
<b>Model Queue Name</b>	Named model queue of the connection.
<b>Max Retries</b>	Maximum number of subsequent connection retry attempts.
<b>Retry Interval</b>	Minimum interval (in seconds) between connection retry attempts.
<b>Wait Interval</b>	Wait interval (in seconds) between attempts to create a connection.
<b>conn</b>	The name of the connection.
<b>time_stamp</b>	The data and time of the last data update.

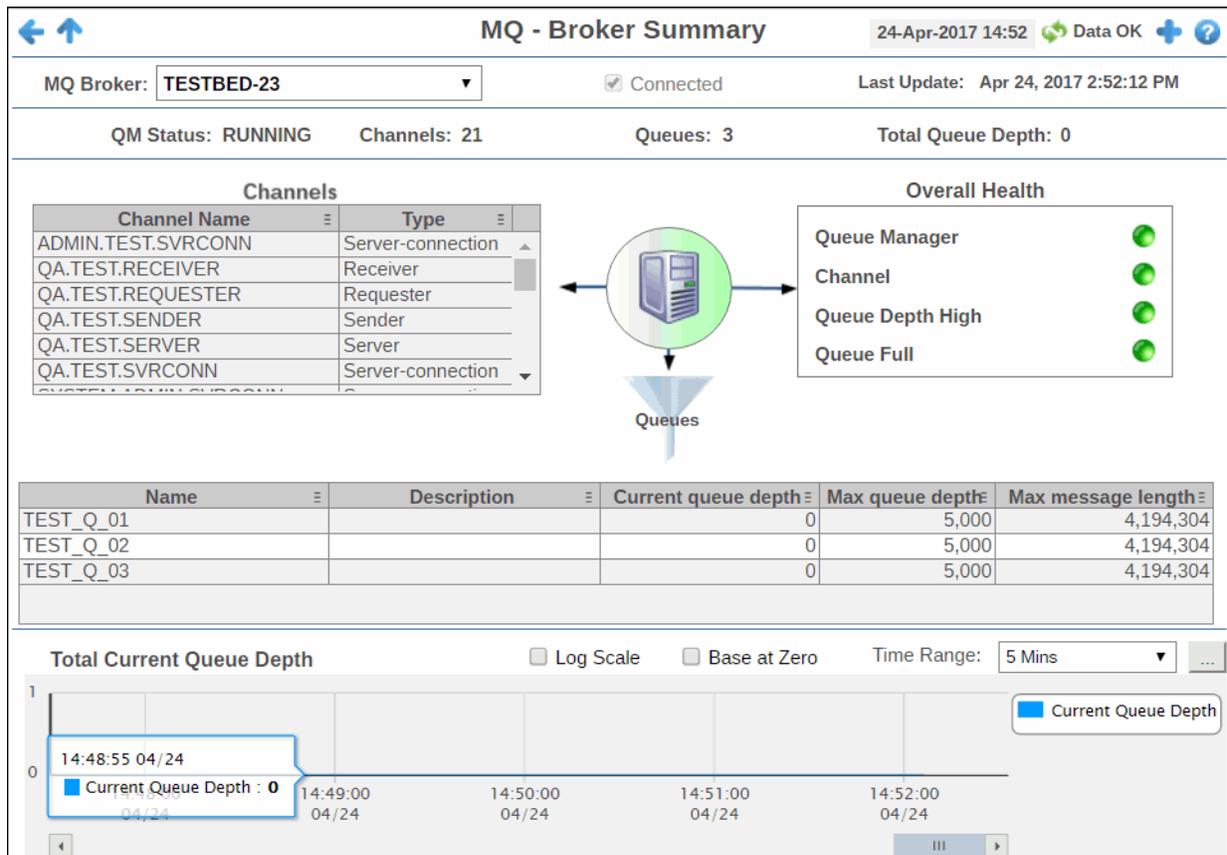
## Single Broker Summary

View detailed utilization metrics and health for a single MQ Broker, such as Channel and Queue health and the number of channels per broker.

Use this display to investigate the performance and health of a broker.

Choose a broker from the drop-down menu. Click a row in the **Channels** table to investigate in the "[All Channels Table](#)" display. Click a row in the **Queues** table. to investigate in the "[All Queues Table](#)" display.

The trend graph traces the total **Current Queue Depth** for the selected broker. .



**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Filter By:**

**MQ Broker** Choose a broker to display.

**Fields and Data**

All values describe the selected broker except where noted.

**Connected** When checked, the broker is connected.

**Last Update** The data and time of the last data update.

**QM Status** The queue manager status. For example:  
**RUNNING** - The QM is operating properly.

**Channels** The current number of channels on the broker.

<b>Queues</b>	The current number of queues on the broker.
<b>Total Queue Depth</b>	The total depth of all queues combined.
<b>Channels</b>	Lists the channels on the broker.  <b>Channel Name:</b> The name of the channel.  <b>Type:</b> The type of channel (for example, Receiver, Cluster-sender, etc.)
<b>Overall Health</b>	The current alert severity for the broker <b>Queue Manager, Channel, Queue Depth High</b> and <b>Queue Full</b> : <ul style="list-style-type: none"> <li>● Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li>● Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li>● Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Queues</b>	Each row is a different queue on the broker. Column values describe the queue.  <b>Name:</b> The name of the queue.  <b>Description:</b> A textual description of the queue.  <b>Current queue depth:</b> The current queue depth.  <b>Max queue depth:</b> The maximum queue depth.  <b>Max message length:</b> The maximum message length in the queue.

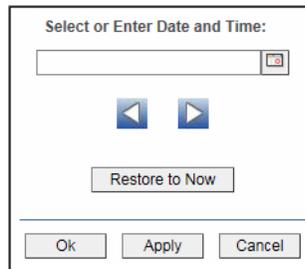
**Trend Graph**

Traces the **Queue Depth** for the selected broker.

<b>Log Scale</b>	Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.
------------------	--

**Base at Zero** Select to use zero as the Y axis minimum for all graph traces.

**Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Channels View

See performance and utilization metrics for all of your IBM MQ Brokers.

Displays in this View are:

- **"All Channels Table"**: This display presents a high-level perspective of utilization metrics for each IBM MQ Broker.
- **"Single Channel Summary"**: This display presents detailed performance metrics for each channel
- **"Single Channel Detail"**: This display presents additional configuration metrics for a single channel.

### All Channels Table

View detailed utilization metrics and parameter settings for MQ Channels on a broker. Metrics include total and delta counts for buffers received/sent. Parameter settings such as MQIACH\_KEEP\_ALIVE\_INTERVAL, MQIACH\_HDR\_COMPRESSION and MQIACH\_MAX\_MSG\_LENGTH are shown.

Each table row is a different channel. Inactive channels are shown in dark red, active channels are shown in green.

Use this display to quickly identify channels with performance issues and confirm channel configurations.

**Note:** This display contains vendor data. Refer to vendor documentation for details.

**MQ - All Channels Detail** 24-Apr-2017 17:29 Data OK

MQ Broker: TESTBED-23

**MQ Channels for Selected Broker**

Channel Name	Type	Alert	Buffers received	Delta	Buffers sent	Delta	Bytes received
ADMIN.TEST.SVRCONN	Server-connection		2,087,011	118	2,087,010	118	994,
QA.TEST.RECEIVER	Receiver		0	0	0	0	
QA.TEST.REQUESTER	Requester		0	0	0	0	
QA.TEST.SENDER	Sender		0	0	0	0	
QA.TEST.SERVER	Server		0	0	0	0	
QA.TEST.SVRCONN	Server-connection		0	0	0	0	
SYSTEM.ADMIN.SVRCONN	Server-connection		0	0	0	0	
SYSTEM.AUTO.RECEIVER	Receiver		0	0	0	0	
SYSTEM.AUTO.SVRCONN	Server-connection		0	0	0	0	
SYSTEM.DEF.CLNTCONN	Client connection		0	0	0	0	
SYSTEM.DEF.CLUSRCVR	Cluster-receiver		0	0	0	0	
SYSTEM.DEF.CLUSDR	Cluster-sender		0	0	0	0	
SYSTEM.DEF.RECEIVER	Receiver		0	0	0	0	
SYSTEM.DEF.REQUESTER	Requester		0	0	0	0	
SYSTEM.DEF.SENDER	Sender		0	0	0	0	
SYSTEM.DEF.SERVER	Server		0	0	0	0	
SYSTEM.DEF.SVRCONN	Server-connection		0	0	0	0	
SYSTEM.FAKE.SVRCONN	Server-connection		0	0	0	0	
WALK.TEST.REQUESTER	Requester		0	0	0	0	
WALKER.ADMIN.SVRCONN	Server-connection		0	0	0	0	
WALKER.TEST.RECEIVER	Receiver		0	0	0	0	

**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu Table open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**MQ Broker** Choose a broker to display.

**MQ Channels for Selected Broker Table**

Each table row is a different channel on the selected broker. Column values describe the channel.

**Channel Name** The name of the channel.

**Type** The type of channel.

<b>Alert</b>	The current alert severity:  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
<b>Buffers Received</b>	The number of buffers received.
<b>Delta</b>	The number of buffers received since the last data update.
<b>Buffers Sent</b>	The number of buffers sent.
<b>Delta</b>	The number of buffers sent since the last data update.
<b>Bytes Received</b>	The number of bytes received.
<b>Delta</b>	The number of bytes received since the last data update.
<b>Bytes Sent</b>	The number of bytes sent.
<b>Delta</b>	The number of bytes sent since the last data update.
<b>Batches Completed</b>	The number of batches completed.
<b>Delta</b>	The number of batches completed since the last data update.
<b>Description</b>	A textual description of the channel.

### Vendor Data

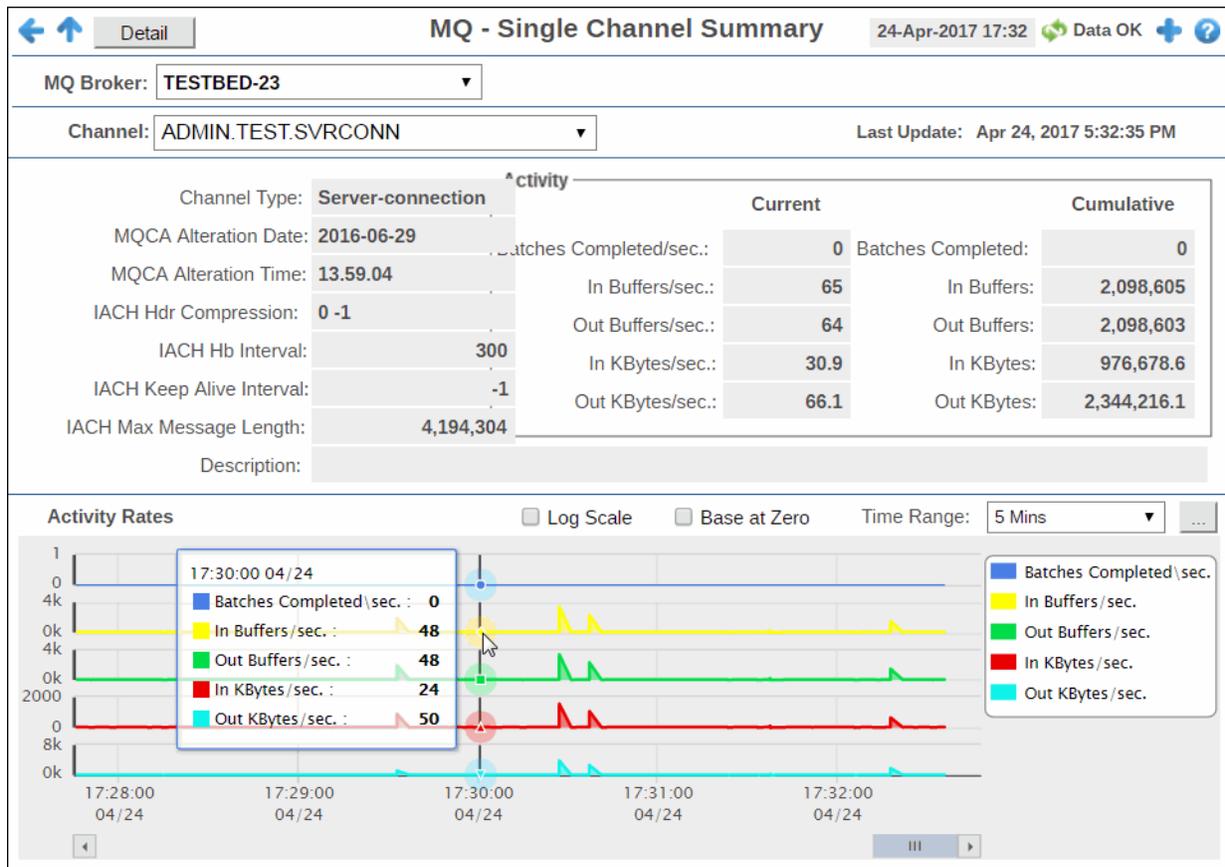
This display contains vendor data. Refer to vendor documentation for details.

<b>Max message length</b>	The maximum length of messages on the channel.
<b>Status</b>	The channel status.
<b>Transmission queue name</b>	The name of the queue that transmits for the channel.
<b>Connection Name</b>	The name of the connection.
<b>time_stamp</b>	The date and time of the last data update.
<b>conn</b>	The name of the connection.

### Single Channel Summary

View current and historical activity rates and transmission settings for a single MQ channel. Metrics include buffers received/sent per second and batches completed. Parameter settings such as MQCA Alteration date, IACH Keep Alive Interval and Max Message Length are shown. Use this display to check the health of a channel and its configuration.

**Note:** This display contains vendor data. Refer to vendor documentation for details.



**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu, Table open commonly accessed displays.
- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

- MQ Broker** Choose a broker to display.
- Channel** Choose a channel to display.
- Last Update** The date and time of the last data update.
- Channel Type** The type of channel.

**Vendor Data**

This display contains vendor data. Refer to vendor documentation for details.

**Activity Rates Trend Graph**

Values describe the selected channel.

	<b>Current</b>	<b>Cumulative</b>
<b>Batches Completed/sec</b>	The current number of batches completed per second.	The total number of batches completed since the channel started.
<b>In Buffers/sec</b>	The current number of buffers received per second.	The total number of buffers received since the channel started.
<b>Out Buffers/sec</b>	The current number of buffers sent per second.	The total number of buffers sent since the channel started.
<b>In KBytes/sec</b>	The current number of kilobytes received per second.	The total number of kilobytes received since the channel started.
<b>Out KBytes/sec</b>	The current number of kilobytes sent per second.	The total number of kilobytes sent since the channel started.

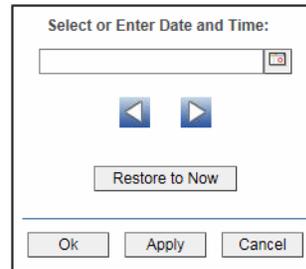
**Activity Rates**

Traces the following for the selected channel:

- **Batches Completed/sec:** The current number batches completed per second.
- **In Buffers/sec:** The current number of buffers received per second.
- **Out Buffers/sec:** The current number of buffers sent per second.
- **In KBytes/sec:** The current number of kilobytes received per second.
- **Out KBytes/sec:** The current number of kilobytes sent per second.

<b>Log Scale</b>	Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.
------------------	--

- Base at Zero** Select to use zero as the Y axis minimum for all graph traces.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

### Single Channel Detail

Get detailed transmission specifications and settings for a single MQ channel. Parameter settings such as CACH Rcv Exit User, IACH Msg Compression and CACH SSL Cipher Spec, as well as DataMQCA Alteration date, IACH Keep Alive Interval and Max Message Length are shown.

**Note:** This display contains vendor data. Refer to vendor documentation for details.

Summary
MQ - Single Channel Detail
24-Apr-2017 17:34 Data OK

MQ Broker: TESTBED-23

Channel: ADMIN.TEST.SVRCONN Last Update: Apr 24, 2017 5:34:14 PM

Channel Type: <span style="border: 1px solid gray; padding: 2px;">Server-connection</span>	MQCA Alteration Date: <span style="border: 1px solid gray; padding: 2px;">2016-06-29</span>
	MQCA Alteration Time: <span style="border: 1px solid gray; padding: 2px;">13.59.04</span>
Description: <span style="border: 1px solid gray; padding: 2px;"> </span>	
IACH Hdr Compression: <span style="border: 1px solid gray; padding: 2px;">0 -1</span>	
IACH Msg Compression: <span style="border: 1px solid gray; padding: 2px;">0 -1 -1 -1 -1 -1 -1 -1 -1 -1 -1 -1</span>	
IACH Hb Interval: <span style="border: 1px solid gray; padding: 2px;">300</span>	IACH MCA User Id: <span style="border: 1px solid gray; padding: 2px;">m</span>
IACH Keep Alive Interval: <span style="border: 1px solid gray; padding: 2px;">-1</span>	IA Monitoring Channel: <span style="border: 1px solid gray; padding: 2px;">-3</span>
IACH Max Message Length: <span style="border: 1px solid gray; padding: 2px;">4,194,304</span>	
CACH Rcv Exit User Data: <span style="border: 1px solid gray; padding: 2px;"> </span>	CACH Rcv Exit Name: <span style="border: 1px solid gray; padding: 2px;"> </span>
CACH Sec Exit User Data: <span style="border: 1px solid gray; padding: 2px;"> </span>	CACH Sec Exit Name: <span style="border: 1px solid gray; padding: 2px;"> </span>
CACH Send Exit User Data: <span style="border: 1px solid gray; padding: 2px;"> </span>	CACH Send Exit Name: <span style="border: 1px solid gray; padding: 2px;"> </span>
IACH SSL Client Auth: <span style="border: 1px solid gray; padding: 2px;">0</span>	CACH SSL Cipher Spec: <span style="border: 1px solid gray; padding: 2px;"> </span>
IACH Xmit Protocol Type: <span style="border: 1px solid gray; padding: 2px;">2</span>	CACH SSL Peer Name: <span style="border: 1px solid gray; padding: 2px;"> </span>

**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu, Table open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

- MQ Broker** Choose a broker to display.
- Channel** Choose a channel to display.
- Last Update** The date and time of the last data update.
- Channel Type** The type of channel. For example, **Server-connection**.

<b>MQCA Alteration Date</b>	The date the MQ CA was last modified.
<b>MQCA Alteration Time</b>	The time the MQ CA was last modified.
<b>Description</b>	The description of the channel definition.
<b>IACH Hdr Compression</b>	The ACH header data compression techniques supported by the channel.
<b>IACH Msg Compression</b>	The ACH message data compression techniques supported by the channel.
<b>IACH Hb Interval</b>	The ACH heartbeat interval setting.
<b>IACH Keep Alive Interval</b>	The ACH keep alive interval setting (the timeout value for the channel).
<b>IACH Max Message Length</b>	The ACH maximum message length setting.
<b>CACH Rcv Exit User Data</b>	The user data that is passed to the receive exit.
<b>CACH Sec Exit User Data</b>	The user data that is passed to the security exit.
<b>CACH Send Exit User Data</b>	The user data that is passed to the send exit.
<b>IACH SSL Client Auth</b>	Denotes whether the channel needs to receive and authenticate an SSL certificate from an SSL client.
<b>IACH Xmit Protocol Type</b>	The transport type used.
<b>IACH MCA User Id</b>	The user ID used by the MCA when attempting to initiate a secure SNA session with a remote MCA.
<b>IA Monitoring Channel</b>	Denotes the attribute used to control the collection of online monitoring data.
<b>CACH Rcv Exit Name</b>	Denotes the name of the user exit program that was run by the channel receive user exit.
<b>CACH Sec Exit Name</b>	Denotes the name of the exit program that was run by the channel security exit.
<b>CACH SSL Cipher Spec</b>	Denotes the single CipherSpec for a TLS or SSL connection.
<b>CACH SSL Peer Name</b>	The Distinguished Name (DN) of the certificate from the peer queue manager or client at the other end of a IBM WebSphere MQ channel.

## Queue Managers

See performance and utilization metrics for all of your IBM MQ queue managers.

Displays in this View are:

- [“All Queue Managers Table”](#): This display presents a high-level perspective of utilization metrics for each IBM MQ queue managers.

### All Queue Managers Table

View detailed utilization metrics and parameter settings for all queue managers on a broker. Metrics include Connection Count and Max Message Length. Parameter settings such as Command Level are shown.

Each table row is a different queue manager. Use this display to quickly identify queue managers with performance issues and confirm configurations.

**Note:** This display contains vendor data. Refer to vendor documentation for details.

Name	Status	Command Level	Connection count	Dead letter queue	Description
UNSECURE	RUNNING	750	28		

**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**MQ Broker** Choose a broker to display.

### MQ Queues Managers for Selected Broker Table

Each table row is a different queue manager on the selected broker. Column values describe the queue.

<b>Name</b>	The name of the queue manager.
<b>Status</b>	The queue manager status (for example, Running).
<b>Command Level</b>	The command level.
<b>Connection Count</b>	The number of connections on the queue manager.
<b>Dead Letter Queue</b>	The number of undelivered messages in the dead letter queue.
<b>Description</b>	A textual description of the queue manager.
<b>Max Message Length</b>	The maximum message length sent or received by the queue manager.
<b>Max Priority</b>	The queue manager rank in priority.
<b>Platform</b>	The queue manager platform type.
<b>Host</b>	The host name.
<b>time_stamp</b>	The date and time of the last data update.
<b>Connection</b>	The connection name.

## Queues

See performance and utilization metrics for all of your IBM MQ queues.

Displays in this View are:

- ["All Queues Table"](#): This display lists all IBM MQ queues with detailed performance metrics and configuration information.
- ["Single Queue Summary"](#): This display presents detailed performance metrics and configuration information for a single IBM MQ queue.

### All Queues Table

View detailed utilization metrics, alert status and settings for IBM MQ queues on a broker. Metrics include total inputs/outputs, depth and Get Messages. Settings such as queue type and default priority are shown.

Choose a broker from the drop-down menu. Each table row is a different queue. Inactive queues are shown in dark red, active queues are shown in green. Investigate by clicking a row to see queue details in the ["Single Queue Summary"](#) display.

Use this display to quickly identify queues with performance or capacity issues and confirm queue configurations.

**Note:** This display contains vendor data. Refer to vendor documentation for details.

Name	Queue Manager	Queue type	Status	Alert	Outputs	Inputs	Dep
TEST_Q_01	UNSECURE	LOCAL	ONLINE	●	0	0	
TEST_Q_02	UNSECURE	LOCAL	ONLINE	●	0	0	
TEST_Q_03	UNSECURE	LOCAL	ONLINE	●	0	0	

**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**MQ Broker** Choose a broker to display.

**MQ Queues for Selected Broker Table**

Each table row is a different queue on the selected broker. Column values describe the queue.

- Name** The name of the queue.
- Queue Manager** The name of the queue manager.
- Queue type** The type of queue.
- Status** The queue status (for example, Online).

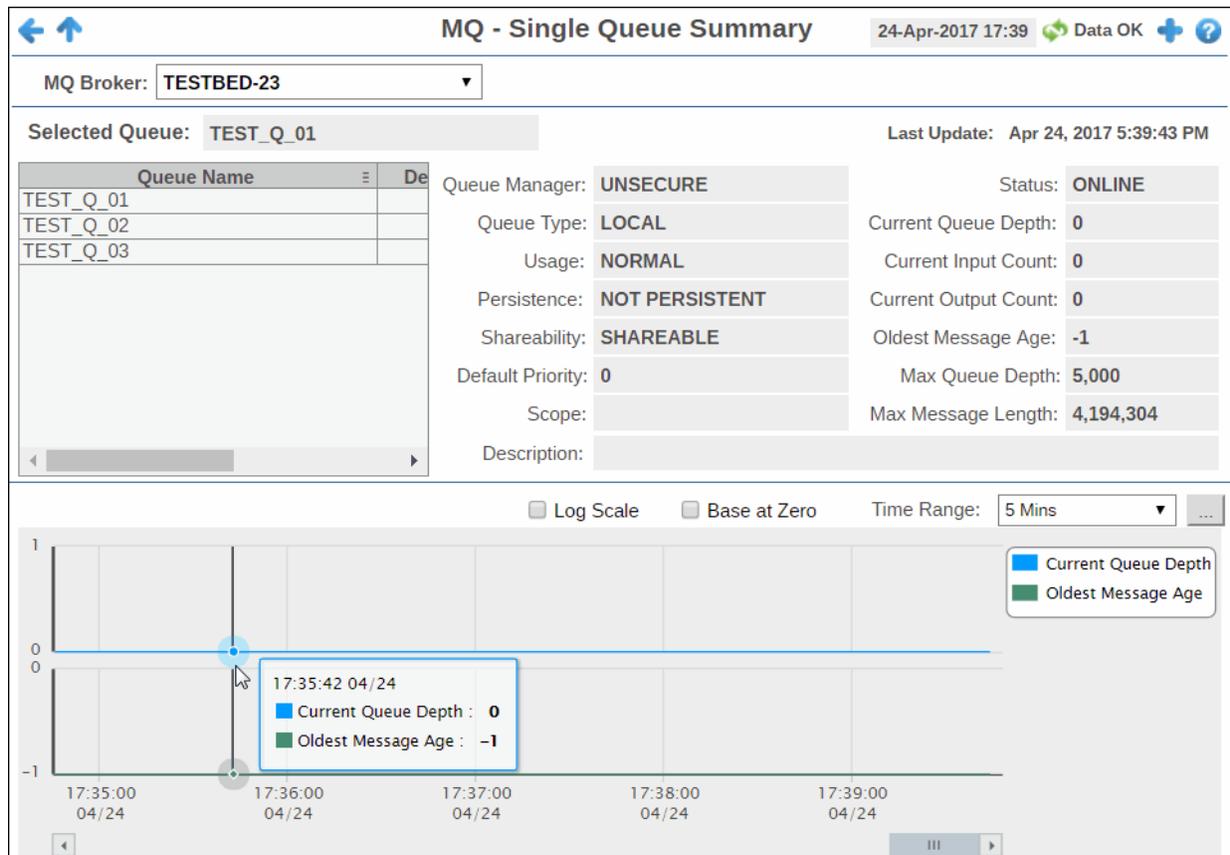
<b>Alert</b>	The current alert severity:  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
<b>Outputs</b>	The number of outgoing transactions.
<b>Inputs</b>	The number of incoming transactions.
<b>Depth</b>	The current queue depth.
<b>Max Depth</b>	The maximum number of messages allowed on the queue at any one time.
<b>Persistence</b>	Denotes whether or not the queue manager is persistent ( <b>PERSISTENT/NOT PERSISTENT</b> ).
<b>Description</b>	The description of the purpose of the queue.
<b>Max Message Length</b>	The maximum length of messages.
<b>Host</b>	The IP address of the host.
<b>Default Priority</b>	The default priority value for messages placed on the queue.
<b>Get Messages</b>	Denotes whether or not the queue enabled to get messages ( <b>GET ALLOWED/GET NOT ALLOWED</b> ).
<b>Put Messages</b>	Denotes whether or not the queue enabled to put messages ( <b>PUT ALLOWED/PUT NOT ALLOWED</b> ).
<b>Scope</b>	The defined scope setting for the queue.
<b>Shareability</b>	Denotes whether or not the queue is shareable ( <b>SHAREABLE/ NOT SHAREABLE</b> ).
<b>Usage</b>	The queue usage type ( <b>NORMAL/TRANSMISSION</b> ).
<b>Connection</b>	The name of the queue connection.
<b>Expired</b>	When checked, this connection is expired due to inactivity.
<b>time_stamp</b>	The date and time of the last data update.

## Single Queue Summary

View detailed performance metrics and settings for a single IBM MQ queue. Metrics include input/output counts, queue depth and max message length. Settings such as persistence, shareability, and scope are shown. Trend graph traces current queue depth and the oldest message age for the selected queue.

Choose a broker from the drop-down menu, then choose a queue from the **Queue** table. Use this display to check the health of a queue and its configuration.

**Note:** This display contains vendor data. Refer to vendor documentation for details.



**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

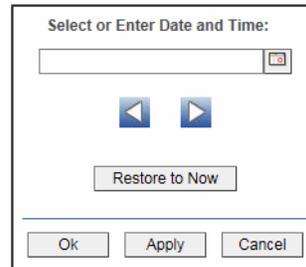
**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

- MQ Broker** Choose a broker to display.
- Selected Queue** Choose a queue from the table to populate the display. The name of the queue selected in the **Queue** table (below).
- Last Update** The data and time of the last data update.
- Queue Table** Choose a queue to populate the display.
  - Queue Name** The queue name.
  - Depth** The current queue depth.
  - Status** The current queue status.

<b>Inputs</b>	The current number of incoming transactions for the queue.
<b>Outputs</b>	The current number of outgoing transactions for the queue.
<b>Queue Manager</b>	The name of the queue manager for the queue.
<b>Queue Type</b>	The type of queue.
<b>Usage</b>	The queue usage type ( <b>NORMAL/TRANSMISSION</b> ).
<b>Persistence</b>	Denotes whether or not the queue manager is persistent ( <b>PERSISTENT/NOT PERSISTENT</b> ).
<b>Shareability</b>	Denotes whether or not the queue manager is shareable ( <b>SHAREABLE/NOT SHAREABLE</b> ).
<b>Default Priority</b>	The default priority setting on the queue manager.
<b>Scope</b>	The defined scope for the queue.
<b>Description</b>	The description of the queue.
<b>Status</b>	The status of the queue.
<b>Current Queue Depth</b>	The current depth of the queue.
<b>Current Input Count</b>	The number of incoming transactions.
<b>Current Output Count</b>	The number of outgoing transactions.
<b>Oldest Message Age</b>	The age of the oldest message in the queue.
<b>Max Queue Depth</b>	The maximum number of messages allowed on the queue.
<b>Max Message Length</b>	The maximum message length on the queue.
<b>Trend Graph</b>	Traces the following for the selected queue: <ul style="list-style-type: none"> <li>• <b>Current Queue Depth</b>: The current depth of the queue.</li> <li>• <b>Oldest Message Age</b>: The age of the oldest message in the queue.</li> </ul>
<b>Log Scale</b>	Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.

- Base at Zero** Select to use zero as the Y axis minimum for all graph traces.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

---

## IBM MQ - HTML

The IBM MQ HTML displays provide extensive visibility into the health and performance of IBM MQ. The HTML version features an ["IBM MQ Overview Display - HTML"](#) (pictured below).

**Note:** This document assumes familiarity with IBM Websphere MQ. For details about IBM MQ, refer to vendor documentation.

The following IBM MQ Views (and their associated displays) can be found under **Components** tab > **Middleware**> **IBM MQ**:

- ["IBM MQ Brokers View - HTML"](#): The displays in this View present performance and utilization metrics for your IBM MQ Brokers.
- ["IBM MQ Channels View - HTML"](#): The displays in this View present performance and utilization metrics for your IBM MQ Channels.
- ["IBM MQ Queues View - HTML"](#): The displays in this View present performance and utilization metrics for your IBM MQ Queue Managers.

## IBM MQ Overview Display - HTML

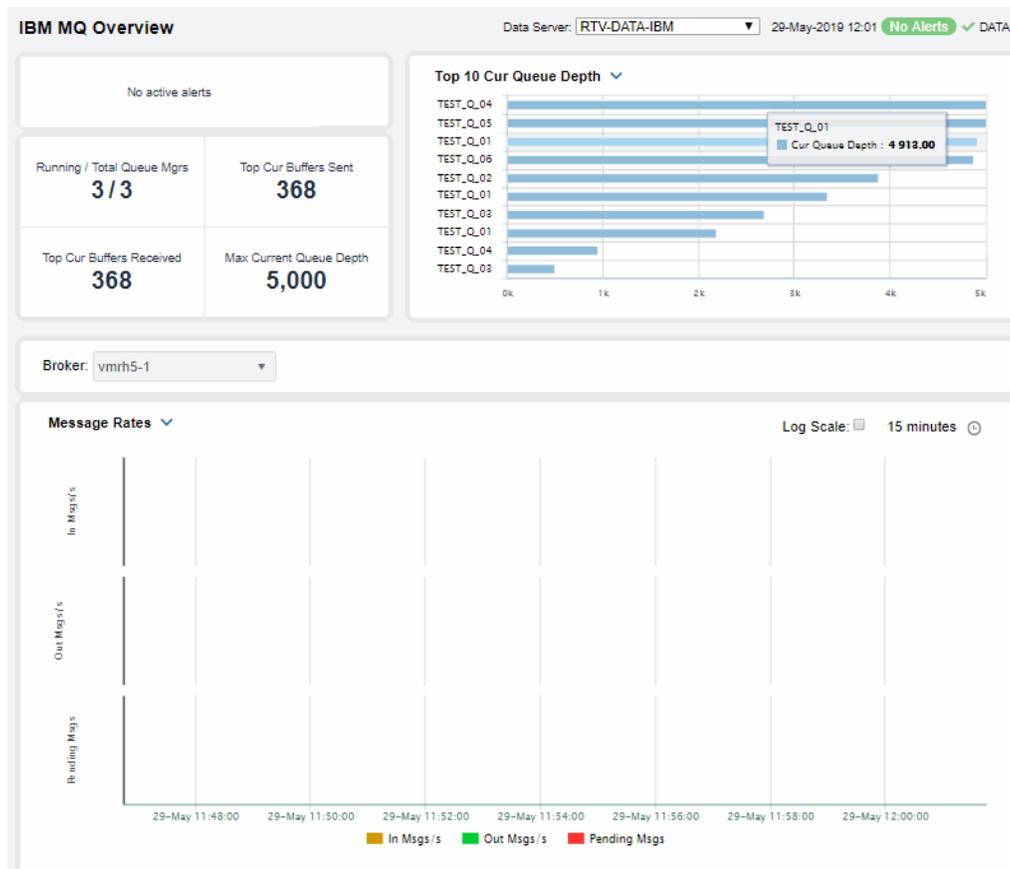
The **IBM MQ Overview** is the top-level display for the IBM MQ Solution Package, which provides a good starting point for immediately getting the status of all IBM MQ objects on your Data Server.

You can select the RTView Data Server for which you want to see data and easily view the current data for that DataServer including:

- The total number of **Active Alerts**, including the total number of critical and warning alerts.
- The number of queues **Running** and the **Total Queue Managers**.
- The number of **Buffers Sent** and **Buffers Received**.
- The **Maximum Current Queue Depth**.
- A bar graphs showing the **Top 10 Current Queue Depth** or the **Oldest Message Age**. Click a bar to drill-down to details in the ["IBM MQ Queue Summary - HTML"](#).

You can hover over each region in the upper half of the Overview to see more detail. You can also drill-down to see more detail (such as in the ["IBM MQ Broker Connections Table - HTML"](#) display) by clicking on each respective region in the Overview.

You can select a **Broker** for the trend graph to trace **Message Rates (In Msgs/second, Out Msgs/second and Pending Msgs)** or **Message Flows**. You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.



## IBM MQ Brokers View - HTML

See performance and utilization metrics for all of your IBM MQ Brokers.

Displays in this View are:

- ["IBM MQ Broker Connections Table - HTML"](#): View a sortable list of utilization metrics for all IBM MQ brokers and compare broker metrics.
- ["IBM MQ Brokers Heatmap - HTML"](#): This display presents heatmap view of IBM MQ brokers and their alert states.
- ["IBM MQ Broker Summary - HTML"](#): This display presents performance metrics for a single IBM MQ Broker, as well as detailed metrics for its channels and queues.

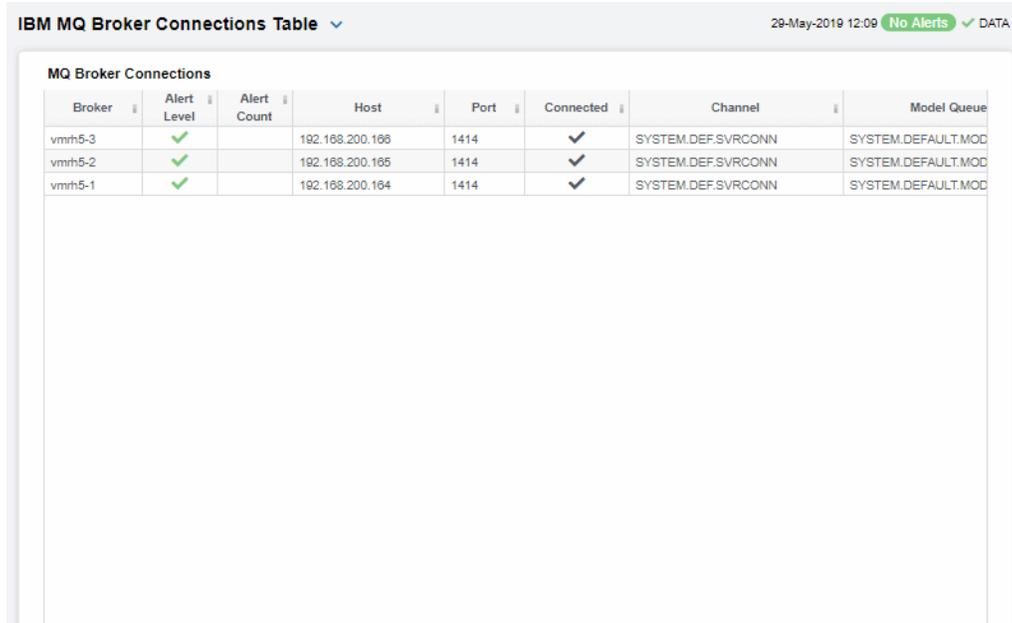
### IBM MQ Broker Connections Table - HTML

Get connection information for all brokers such as **Host** IP address and **Port** number and **Channel Model Queue Name**, and investigate utilization metrics for all brokers such as **Alert Level**, **Alert Count**, **Max Retries** and **Wait Interval**.

Each row in the table contains data for a particular broker. You can search, filter, sort and choose columns to include by clicking a column header icon (to the right of each column label) and selecting **Filter**, **Sort Ascending**, **Sort Descending** or **Columns**.

Or just click any column header to sort and compare the values that interest you. Right-click on a table cell to **Export to Excel** or **Copy Cell Value**.

Investigate a broker, its channels and queues by double-clicking a row which opens the ["IBM MQ Broker Summary - HTML"](#) display.



Broker	Alert Level	Alert Count	Host	Port	Connected	Channel	Model Queue
vmrh5-3	OK		192.168.200.166	1414	✓	SYSTEM.DEF.SVRCONN	SYSTEM.DEFAULT.MOQ
vmrh5-2	OK		192.168.200.165	1414	✓	SYSTEM.DEF.SVRCONN	SYSTEM.DEFAULT.MOQ
vmrh5-1	OK		192.168.200.164	1414	✓	SYSTEM.DEF.SVRCONN	SYSTEM.DEFAULT.MOQ

### IBM MQ Brokers Heatmap - HTML

View current alert status and performance metrics of all IBM MQ brokers. Use the **Metric** drop-down menu to view **Alert Severity**, **Alert Count** and **Current Queue Depth**.

Each rectangle in the heatmap represents a different broker, where the rectangle color indicates the most critical alert state for items associated with that broker. The rectangle size is the same for each broker.

By default, the **Alert Severity** metric is shown. Values range from **0** - **2**, as indicated in the color gradient  bar:

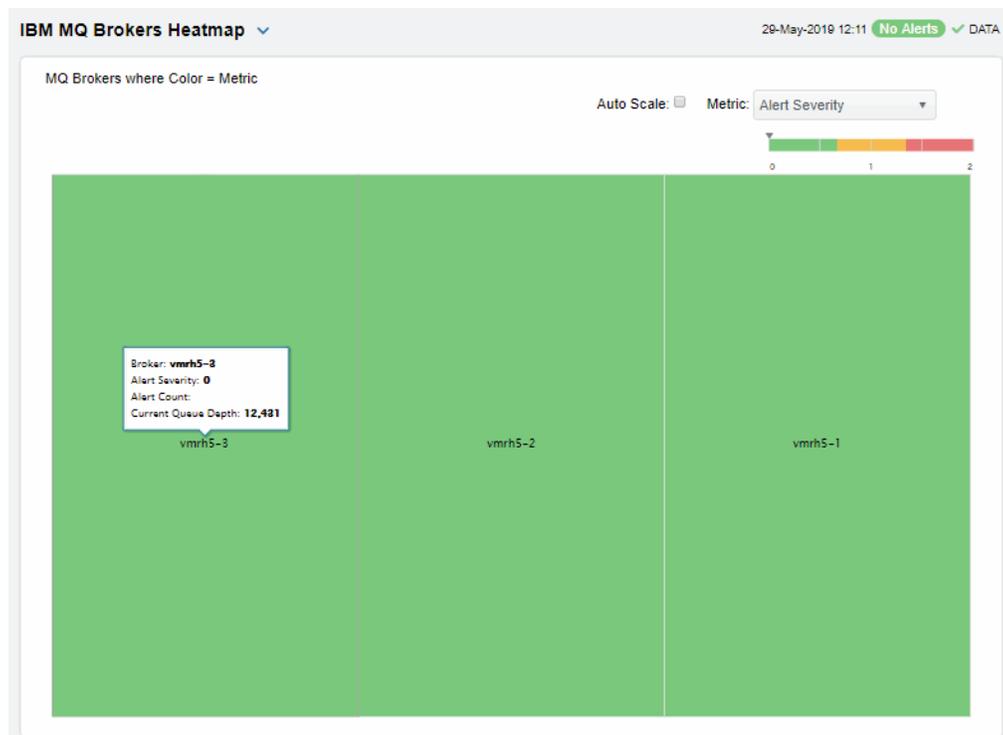
- (2) Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
- (1) Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
- (0) Green indicates that no metrics have exceeded their alert thresholds.

Answer questions such as, Are any queues reaching a state of critical health? Is the load evenly distributed across brokers and queues?

Investigate a broker, its channels and queues by clicking a rectangle which opens the ["IBM MQ Broker Summary - HTML"](#) display.

Mouse-over rectangles to view more details about host performance and status. Toggle between the commonly accessed Table and Heatmap displays by clicking the drop down list on the display title.

You can view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.



## IBM MQ Broker Summary - HTML

Investigate the performance and health of a particular IBM MQ broker, its channels and its queues. Track utilization and performance metrics of channels and queues for a particular IBM MQ broker in a trend graph.

Choose a broker from the drop-down menu. Clicking on the information boxes at the top of the display (such as **Running Queue Mgrs**, **Total Queue Depth** and **Total Open Inputs/Outputs**) takes you to the ["IBM MQ Broker Connections Table - HTML"](#) display, where you can sort and compare the performance values of all brokers.

The trend graph traces the **Current Queue Depth** for the selected broker. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.

The **MQ Channels for Selected Broker** table (the lower portion of the display) contains a row of data for each channel on the broker. Double-click a row to investigate the channel in the ["IBM MQ Channel Summary - HTML"](#) display. You can search, filter, sort and choose columns to include by clicking a column header icon (to the right of each column label) and selecting **Filter**, **Sort Ascending**, **Sort Descending** or **Columns**.

Or just click any column header to sort and compare the values that interest you. Right-click on a table cell to **Export to Excel** or **Copy Cell Value**.

Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display. You can hover over the trend graph to see the values at a particular time.



## IBM MQ Channels View - HTML

See performance and utilization metrics for all of your IBM MQ Brokers.

Displays in this View are:

- ["IBM MQ Channels Table - HTML"](#): This display presents a high-level perspective of utilization metrics for each IBM MQ Broker.
- ["IBM MQ Channels Heatmap - HTML"](#): This display presents detailed performance metrics for each channel
- ["IBM MQ Channel Summary - HTML"](#): This display presents additional configuration metrics for a single channel.

## IBM MQ Channels Table - HTML

View performance metrics (such as **Alert Level**, **Rcvd Buffers/Sent Buffers** and **Completed Batches**) connection information (such as channel **Type**, **Connection Name** and **Transmission Queue Name**) for all channels on a broker. Also included are parameter settings such as **MQIACH\_KEEP\_ALIVE\_INTERVAL**, **MQIACH\_HDR\_COMPRESSION** and **MQIACH\_MAX\_MSG\_LENGTH** are shown.

Use this display to quickly identify channels with performance issues and confirm channel configurations.

Each row in the table contains data for a particular channel. You can search, filter, sort and choose columns to include by clicking a column header icon (to the right of each column label) and selecting **Filter**, **Sort Ascending**, **Sort Descending** or **Columns**.

Or just click any column header to sort and compare the values that interest you. Right-click on a table cell to **Export to Excel** or **Copy Cell Value**.

Investigate a broker, its channels and queues by double-clicking a row which opens the ["IBM MQ Channel Summary - HTML"](#) display.

**Note:** This display contains vendor data. Refer to vendor documentation for details.

Channel Name	Type	Alert Level	Alert Count	Rcvd Buffers	Current Rcvd Buffers	Sent Buffers	Current Sent Buffers	Rcvd By
		✓		0	0	0	0	
SYSTEM.DEF.SVRCONN	Server-connection	✓		112,857	0	112,855	0	99.2
SYSTEM.DEF.SVRCONN	Server-connection	✓		111,541	341	111,539	342	98.1
SYSTEM.DEF.SVRCONN	Server-connection	✓		227,465	0	227,463	0	14.5

## IBM MQ Channels Heatmap - HTML

View current alert status and performance metrics for all channels on a particular IBM MQ broker, or all channels on **All** brokers.

Use the **Metric** drop-down menu to view **Alert Severity**, **Alert Count**, **Current Buffers Received**, **Current Buffers Sent**, **Current Bytes Received** or **Current Bytes Sent**.

Each rectangle in the heatmap represents a different channel, where the rectangle color indicates the most critical alert state for items associated with that broker. The rectangle size is the same for each channel.

By default, the **Alert Severity** metric is shown. Values range from **0** - **2**, as indicated in the color gradient  bar:

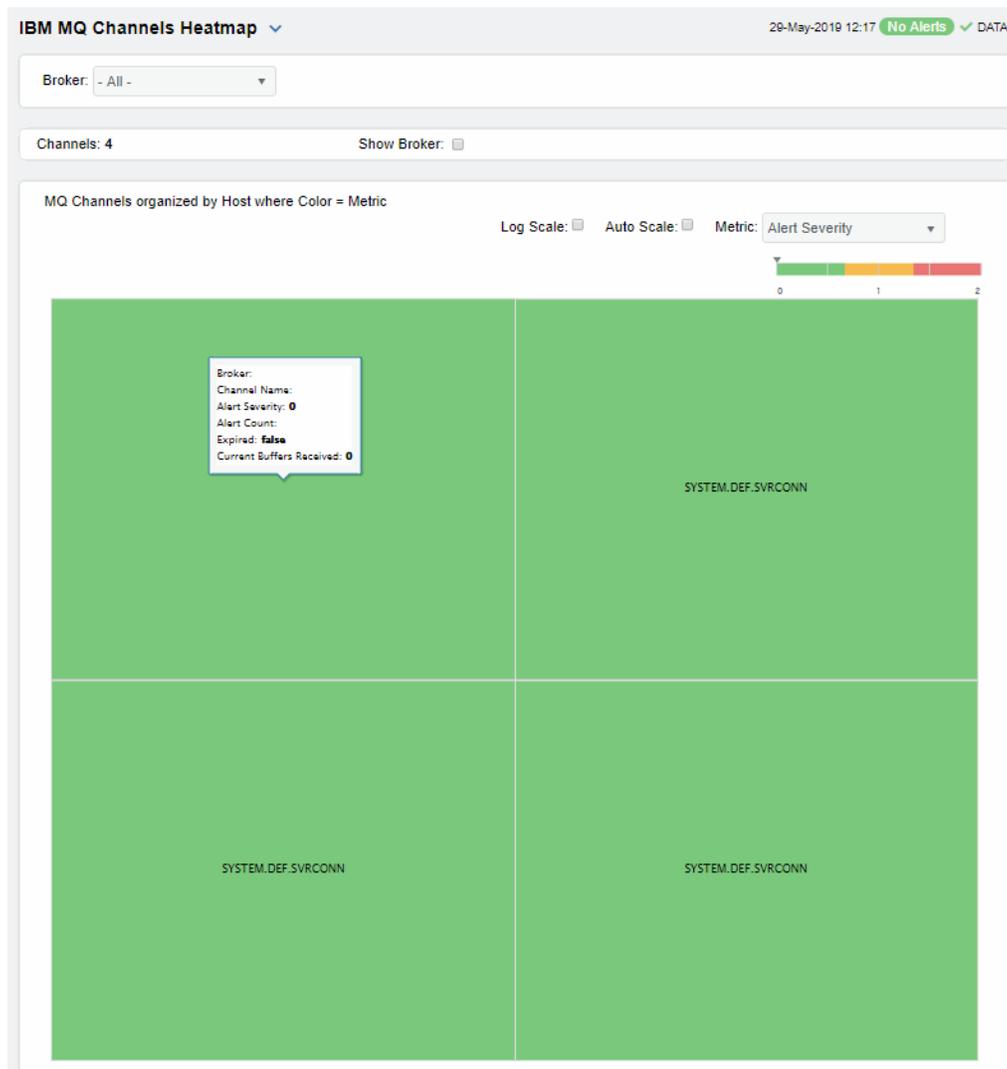
- **(2)** Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
- **(1)** Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
- **(0)** Green indicates that no metrics have exceeded their alert thresholds.

Answer questions such as, Are any channels reaching a state of critical health? Is the load evenly distributed across channels?

Investigate a channel by clicking a rectangle which opens the ["IBM MQ Channel Summary - HTML"](#) display.

Mouse-over rectangles to view more details about host performance and status. Toggle between the commonly accessed Table and Heatmap displays by clicking the drop down list on the display title.

You can view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.



## IBM MQ Channel Summary - HTML

Investigate the performance and health of a particular IBM MQ channel. View detailed transmission metrics and settings for a single IBM MQ channel, such as **IACH Hdr Compression** and **IACH Max Message Length**, as well as **MQCA Alteration Time**, **IACH Keep Alive Interval** and **Completed Batches** are shown.

Track utilization and performance metrics of channels and queues for a particular IBM MQ broker in a trend graph.

Use this display to check the health of a channel and its configuration.

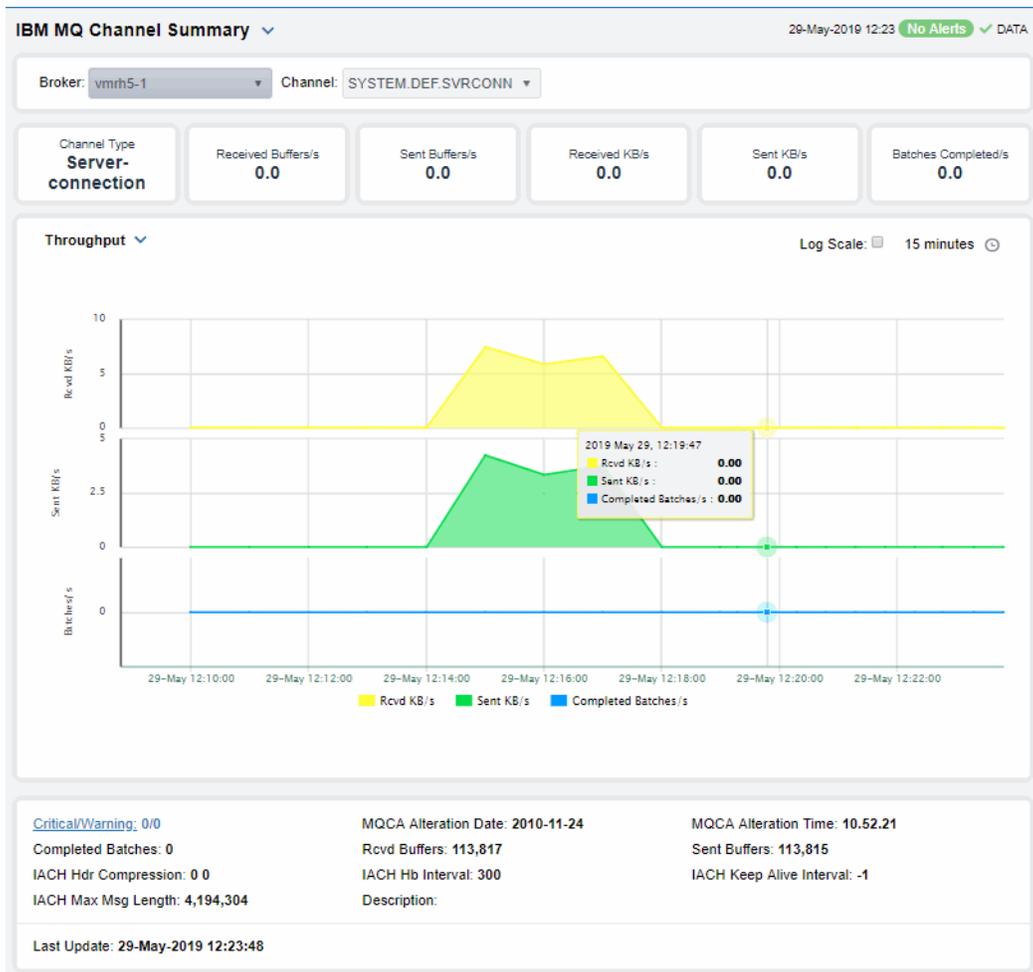
Choose a broker and a channel from the drop-down menus. Clicking on the information boxes at the top of the display (such as **Received Buffers/second**, **Sent Buffers/second** and **Batches Completed**) takes you to the "IBM MQ Channels Table - HTML" display, where you can sort and compare the performance values of all channels.

The trend graph traces the **Throughput** and **Buffer Flow** rates for the selected channel. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.

Or just click any column header to sort and compare the values that interest you. Right-click on a table cell to **Export to Excel** or **Copy Cell Value**.

Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display. You can hover over the trend graph to see the values at a particular time.

**Note:** This display contains vendor data. Refer to vendor documentation for details.



## IBM MQ Queues View - HTML

See performance and utilization metrics for all of your IBM MQ queue managers.

Displays in this View are:

- [“IBM MQ Queue Managers Table - HTML”](#): This display presents a high-level perspective of utilization metrics for each IBM MQ queue managers.
- [“IBM MQ Queues Table - HTML”](#): This display lists all IBM MQ queues with detailed performance metrics and configuration information.
- [“IBM MQ Queues Heatmap - HTML”](#): This display presents a heatmap view of performance metrics and alert levels for one or all brokers.
- [“IBM MQ Queues Heatmap - HTML”](#): This display presents detailed performance metrics and configuration information for a single IBM MQ queue.

## IBM MQ Queue Managers Table - HTML

View detailed utilization metrics and parameter settings for all queue managers on a particular broker or on all brokers.

Each table row contains data for a particular queue manager. Use this display to quickly identify queue managers with performance issues and confirm configurations. Metrics include **Connection Count** and **Max Message Length**. Parameter settings such as **Command Level** are shown.

Use this display to quickly identify brokers with performance issues.

You can search, filter, sort and choose columns to include by clicking a column header icon (to the right of each column label) and selecting **Filter**, **Sort Ascending**, **Sort Descending** or **Columns**.

Or just click any column header to sort and compare the values that interest you. Right-click on a table cell to **Export to Excel** or **Copy Cell Value**.

Investigate a broker, its channels and queues by double-clicking a row which opens the [“IBM MQ Broker Summary - HTML”](#) display.

**Note:** This display contains vendor data. Refer to vendor documentation for details.

IBM MQ Queue Managers Table 29-May-2019 15:43 No Alerts DATA

Broker: - All -

Broker	Name	Alert Level	Alert Count	Status	Command Level	Connection count	Dead letter queue	Descripti
vmrh5-3	QM_0001	✓		RUNNING	701	52	DLQ	
vmrh5-2	QM_0001	✓		RUNNING	701	23		
vmrh5-1	QM_0001	✓		RUNNING	701	28		

## IBM MQ Queues Table - HTML

View performance metrics, alert status and settings for IBM MQ queues on a particular broker or on all brokers. Metrics include total **Get Messages/Put Messages**, **Max Q Depth** and **Current Q Depth**. Settings such as queue **Type**, **Host** IP address and **Persistence** are shown.

Use this display to quickly identify queues with performance issues or capacity issues and confirm configurations.

Each row in the table contains data for a particular queue. You can search, filter, sort and choose columns to include by clicking a column header icon (to the right of each column label) and selecting **Filter**, **Sort Ascending**, **Sort Descending** or **Columns**.

Or just click any column header to sort and compare the values that interest you. Right-click on a table cell to **Export to Excel** or **Copy Cell Value**.

Investigate a queue by double-clicking a row which opens the ["IBM MQ Queue Summary - HTML"](#) display.

**Note:** This display contains vendor data. Refer to vendor documentation for details.

IBM MQ Queues Table 29-May-2019 15:45 No Alerts DATA

Broker: - All -

Queues: 30

Queue Name	Queue Manager	Queue Type	Status	Alert Level	Alert Count	Open Outputs	Open Inputs	Cur Queue Depth	Max Q Depth	Persist
TEST_Q_01	QM_0001	LOCAL	ONLINE	✓		0	0	4,913	5,000	NOT PER
TEST_Q_02	QM_0001	LOCAL	ONLINE	✓		0	0	3,880	5,000	NOT PER
TEST_Q_03	QM_0001	LOCAL	ONLINE	✓		0	0	0	5,000	NOT PER
TEST_Q_04	QM_0001	LOCAL	ONLINE	✓		0	0	947	5,000	NOT PER
TEST_Q_05	QM_0001	LOCAL	ONLINE	✓		2	1	4,873	5,000	NOT PER
TEST_Q_06	QM_0001	LOCAL	ONLINE	✓		0	0	4,867	5,000	NOT PER
TEST_Q_07	QM_0001	LOCAL	ONLINE	✓		0	0	0	5,000	NOT PER
TEST_Q_08	QM_0001	LOCAL	ONLINE	✓		0	0	0	5,000	NOT PER
TEST_Q_09	QM_0001	LOCAL	ONLINE	✓		0	0	0	5,000	NOT PER
TEST_Q_10	QM_0001	LOCAL	ONLINE	✓		0	0	0	5,000	NOT PER
TEST_Q_01	QM_0001	LOCAL	ONLINE	✓		2	2	0	5,000	NOT PER
TEST_Q_02	QM_0001	LOCAL	ONLINE	✓		0	0	0	5,000	NOT PER
TEST_Q_03	QM_0001	LOCAL	ONLINE	✓		0	0	504	5,000	NOT PER
TEST_Q_04	QM_0001	LOCAL	ONLINE	✓		0	0	0	5,000	NOT PER
TEST_Q_05	QM_0001	LOCAL	ONLINE	✓		0	0	433	5,000	NOT PER
TEST_Q_06	QM_0001	LOCAL	ONLINE	✓		0	0	0	5,000	NOT PER
TEST_Q_07	QM_0001	LOCAL	ONLINE	✓		0	0	0	5,000	NOT PER
TEST_Q_08	QM_0001	LOCAL	ONLINE	✓		0	0	0	5,000	NOT PER
TEST_Q_09	QM_0001	LOCAL	ONLINE	✓		0	0	0	5,000	NOT PER
TEST_Q_10	QM_0001	LOCAL	ONLINE	✓		0	0	0	5,000	NOT PER
TEST_Q_01	QM_0001	LOCAL	ONLINE	✓		1	1	23	5,000	NOT PER
TEST_Q_02	QM_0001	LOCAL	ONLINE	✓		0	1	0	5,000	NOT PER
TEST_Q_03	QM_0001	LOCAL	ONLINE	✓		1	1	0	5,000	NOT PER
TEST_Q_04	QM_0001	LOCAL	ONLINE	✓		1	0	5,000	5,000	NOT PER
TEST_Q_05	QM_0001	LOCAL	ONLINE	✓		0	0	0	5,000	NOT PER
TEST_Q_06	QM_0001	LOCAL	ONLINE	✓		2	3	0	5,000	NOT PER
TEST_Q_07	QM_0001	LOCAL	ONLINE	✓		0	0	0	5,000	NOT PER
TEST_Q_08	QM_0001	LOCAL	ONLINE	✓		0	0	0	5,000	NOT PER
TEST_Q_09	QM_0001	LOCAL	ONLINE	✓		0	0	0	5,000	NOT PER
TEST_Q_10	QM_0001	LOCAL	ONLINE	✓		0	0	0	5,000	NOT PER

### IBM MQ Queues Heatmap - HTML

View current alert status and performance metrics for all queues on a particular IBM MQ broker, or all queues on **All** brokers.

Use the **Metric** drop-down menu to view **Alert Severity**, **Alert Count**, **Max Queue Depth**, **Max Message Length**, **Open Inputs** or **Open Outputs**.

Each rectangle in the heatmap represents a different queue, where the rectangle color indicates the most critical alert state for items associated with that broker. The rectangle size is the same for each queue.

By default, the **Alert Severity** metric is shown. Values range from **0** - **2**, as indicated in the color gradient  bar:

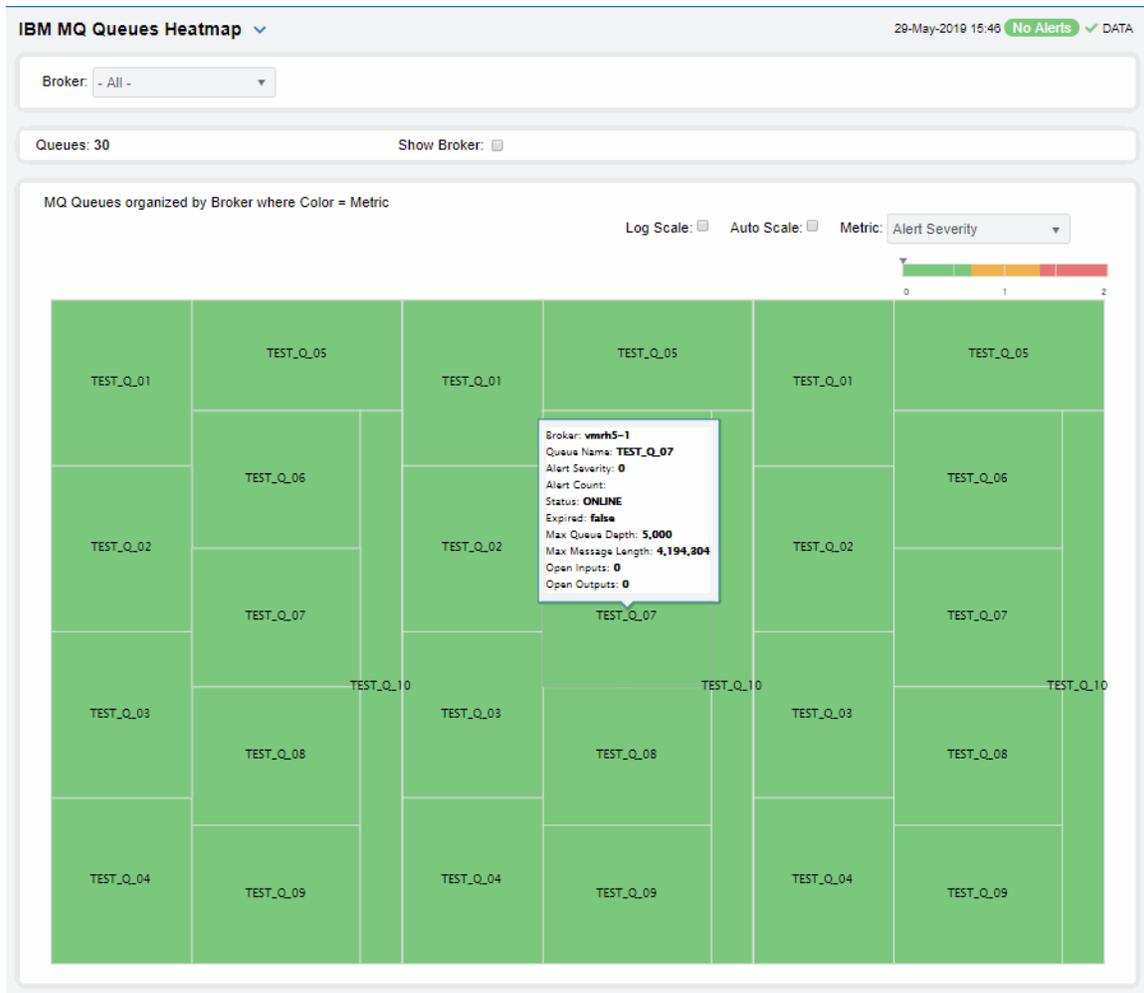
- **(2)** Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
- **(1)** Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
- **(0)** Green indicates that no metrics have exceeded their alert thresholds.

Answer questions such as, Are any queues reaching a state of critical health? Is the load evenly distributed across queues?

Investigate a channel by clicking a rectangle which opens the “[IBM MQ Queue Summary - HTML](#)” display.

Mouse-over rectangles to view more details about host performance and status. Toggle between the commonly accessed Table and Heatmap displays by clicking the drop down list on the display title.

**Note:** This display contains vendor data. Refer to vendor documentation for details.



## IBM MQ Queue Summary - HTML

Investigate the performance and health of a particular IBM MQ queue. View detailed transmission metrics and settings for a single IBM MQ queue, such as **Open Outputs**, **Queue Type** and **Persistence**.

Track utilization and performance metrics of a queue on a particular IBM MQ broker.

Use this display to check the health of a channel and its configuration.

Choose a broker and a queue from the drop-down menus. Clicking on the information boxes at the top of the display (such as **Received Buffers/second**, **Sent Buffers/second** and **Batches Completed**) takes you to the "IBM MQ Channels Table - HTML" display, where you can sort and compare the performance values of all channels.

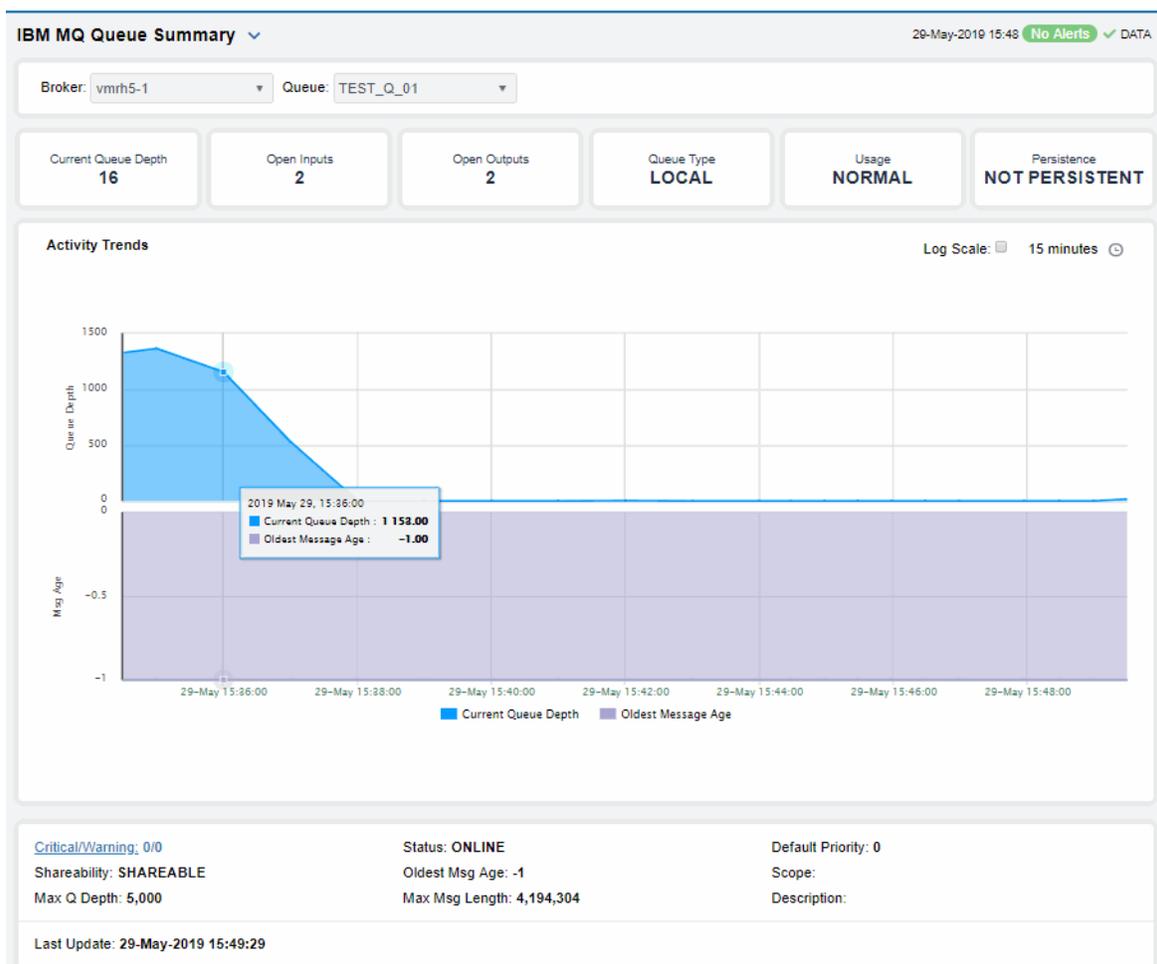
The trend graph traces the **Activity Trends (Current Queue Depth and Oldest Message Age)** rates for the selected queue. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.

Or just click any column header to sort and compare the values that interest you. Right-click on a table cell to **Export to Excel** or **Copy Cell Value**.

Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display. You can hover over the trend graph to see the values at a particular time.

**Shareability, Max Message Length, Status** and **Default Priority** settings for the queue are also shown at the bottom of the display.

**Note:** This display contains vendor data. Refer to vendor documentation for details.



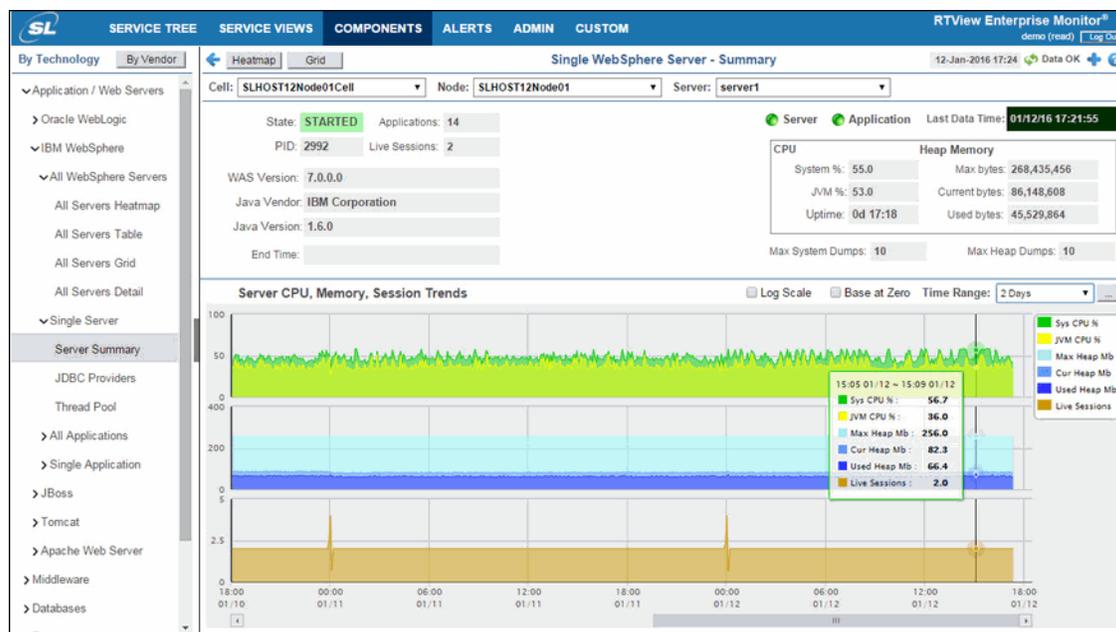
## IBM WebSphere

Gain real-time visibility into the health and performance of WebSphere Application Server and deployed applications. The following WebSphere Monitor Views (and their associated displays) can be found under **Components** tab > **Application/Web Servers**> **WebSphere**:

- **"All WebSphere Servers View"**: The displays in this View present high-level performance and utilization metrics for all of your IBM WebSphere servers.
- **"Single Server View"**: The displays in this View present high-level performance and utilization metrics for a single IBM WebSphere server.
- **"All Applications View"**: The displays in this View present high-level performance and utilization metrics for all of your web application sessions.
- **"Single Application View"**: The displays in this View present high-level performance and utilization metrics for a single web application.

## Overview

RTView Enterprise and the Solution Package for IBM® WebSphere provide out of box performance and availability monitoring for support teams and WebSphere administrators. It enables users to ensure effective resource allocation by providing access to a wide variety of current and historical metrics.



Configuration options enable both consolidated views across the enterprise as well as views configured for specific support teams. As part of an end to end monitoring solution, users can view WebSphere performance in the context of an application or service. This provides visibility into how WebSphere performance is impacting adjacent technologies and the resulting business impact. Typical installations of RTView Enterprise and its solution packages take only a few hours, while developing custom views for a variety of IT and development roles can be achieved in just days.

## Key Features

- Monitor real-time performance for early warning
- Analyze historical performance to differentiate trends and spikes
- Out of the box discovery and monitoring of key metrics and resources
- Ensure effective resource allocation
- Powerful diagnostics and correlations for complex performance analysis
- View WebSphere performance in an application context

## Metrics for WebSphere Server

- System CPU Usage & Process CPU Usage
- Uptime
- Max Memory
- Heap Size & Max Heap Dumps on Disks
- Java Vendor & Version
- Used Memory, Free Memory & Used Memory Percent
- JVM Memory
- Max Heap, Current Heap, Used Heap
- Live Sessions
- JDBC Providers: Open Count, Created, Pool Size, Used Pool, Use Time
- Thread Pools: Pool Size, Active Count, Growable Indication, Inactivity Timeout & Max size
- Server Applications Sessions/Requests: Number of Sessions, Servlets, Total Requests, Current Requests, Avg. Response Time
- JSP Requests, JSP Response Time, Servlet Requests, Servlet Response Time, EJB Method Calls, EJB Response Time
- Component Detail
- Module Detail Totals for Charts & Tables

## End-to-End Context for WebSphere

- Custom flow diagrams help visualize complex applications and WebSphere's place in that architecture
- Provides an Intuitive View of How WebSphere Interacts with other Enterprise PaaS Components
- Designed and Developed for Large Scale, Mission Critical Environments

## All WebSphere Servers View

See performance and utilization metrics for all of your IBM WebSphere servers.

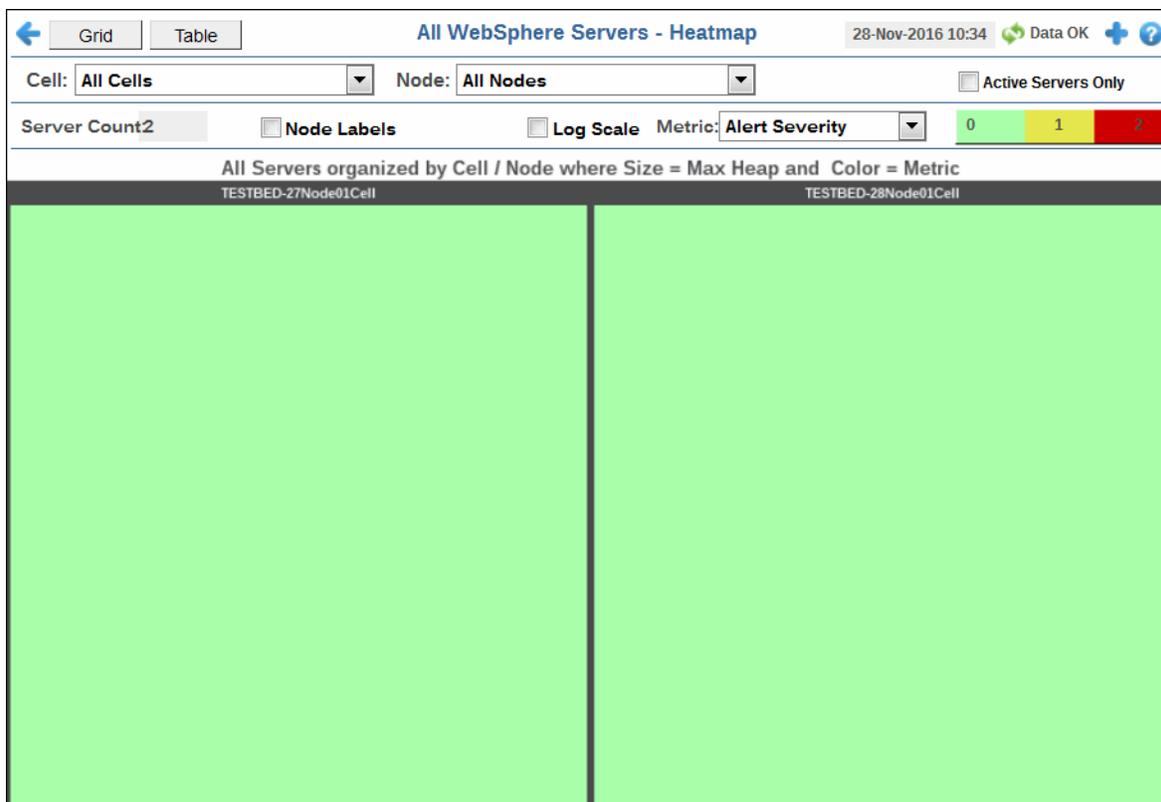
Displays in this View are:

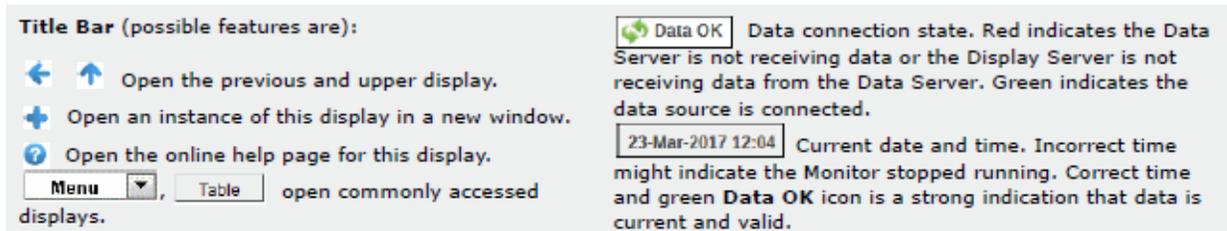
- ["All Servers Heatmap"](#)
- ["All Servers Table"](#)
- ["All Servers Grid"](#)
- ["All Servers Detail"](#)

### All Servers Heatmap

This heatmap shows the current status and utilization metrics for all WebSphere servers. Choose a cell and node from the drop-down menus. Use this display to see metrics for **Alert Count, Live Session Count, WAS CPU %, Host CPU %** and **Memory Used %**. By default, this display shows the heatmap based on the **Alert Severity** metric.

Each rectangle is a different WebSphere server. Use the **Node Labels** check-box  to include or exclude labels in the heatmap, and mouse over a rectangle to see additional metrics for a server. Click a rectangle to drill down to the ["Server Summary"](#) display, which shows additional details about the server.





**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- ,  open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

### Filter By:

- Cell:** Choose a cell or **All Cells** to see metrics for.
- Node:** Choose a node or **All Nodes** to see metrics for.

### Fields and Data:

- Active Servers Only** Choose this check box to only include active servers in the display.
- Server Count** The number of servers in the display.
- Node Labels** Select to include node labels in the display.
- Log Scale** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Metric** Choose a metric to view in the display.
  - Alert Severity** The current alert severity. Values range from **0** - **2**, as indicated in the color gradient bar, where **2** is the highest Alert Severity:
    - Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
    - Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
    - Green indicates that no metrics have exceeded their alert thresholds.
  - Alert Count** The total number of critical and warning unacknowledged alerts. The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.
  - CPU Used%** The percent CPU used. The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **ProcessCpuLoadPercent**. The middle value in the gradient bar indicates the middle value of the range.
 

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

**V Memory Used%**

The percent virtual memory used. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **SystemCpuLoadPercent**. The middle value in the gradient bar indicates the middle value of the range.

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

**Free Memory**

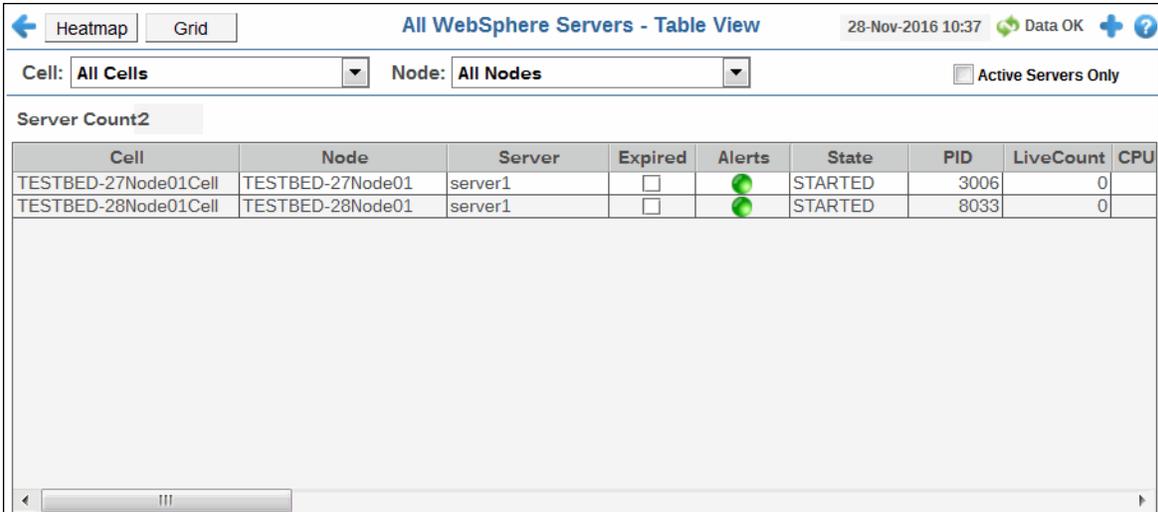
The total amount of available memory. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum amount of available memory. The middle value in the gradient bar indicates the average amount.

The **Auto** flag does not have any impact on this metric.

**All Servers Table**

View WebSphere server utilization details, including memory, CPU and heap size. Choose a cell and node from the drop-down menus. Each row in the table is a different server. The row color for inactive servers is dark red.

Drill-down and investigate by clicking a row in the table to view details for the selected connection in the ["Server Summary"](#) display.



Cell	Node	Server	Expired	Alerts	State	PID	LiveCount	CPU
TESTBED-27Node01Cell	TESTBED-27Node01	server1	<input type="checkbox"/>		STARTED	3006	0	
TESTBED-28Node01Cell	TESTBED-28Node01	server1	<input type="checkbox"/>		STARTED	8033	0	

**Title Bar** (possible features are):

-   Open the previous and upper display.
-  Open an instance of this display in a new window.
-  Open the online help page for this display.
-   open commonly accessed displays.

 **Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

 **23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Filter By:**

**Cell:** Choose a cell or **All Cells** to see metrics for.

**Node:** Choose a node or **All Nodes** to see metrics for.

**Fields and Data**

This display includes:

<b>Active Servers Only</b>	Select to only include active servers in the display.
<b>Server Count</b>	The number of servers in the display.

**Table**

Each table row is a different server. Table column values describe the cell on the server.

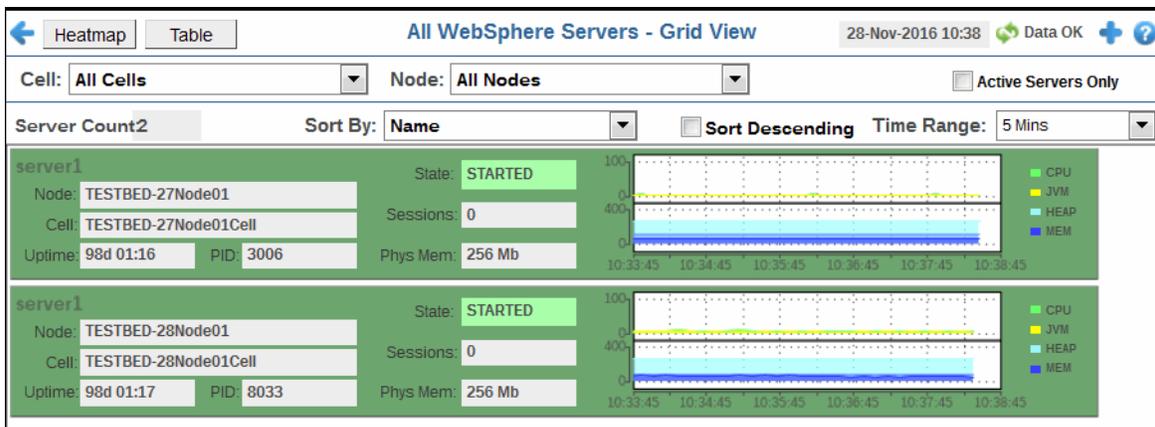
<b>Cell</b>	The name of the cell.
<b>Node</b>	The name of the node.
<b>Server</b>	The name of the server.
<b>Expired</b>	When checked, data has not been received from this host in the specified amount of time. The host will be removed from the Monitor in the specified amount of time. The default setting is <b>60</b> seconds.
<b>Alert Level</b>	The current alert severity.  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
<b>State</b>	The WebSphere server current state: <ul style="list-style-type: none"> <li>• STARTED</li> <li>• STOPPED</li> </ul>
<b>PID</b>	The WebSphere server process identifier.
<b>LiveCount</b>	The total number of currently active sessions.
<b>CPUUsageSinceLastMeasurement</b>	The amount of CPU usage, in megabytes, since the last data update.
<b>CPUUsageSinceServerStarted</b>	The amount of CPU usage, in megabytes, since the server was started.
<b>ProcessCpuUsage</b>	The amount of process CPU usage, in megabytes, since the server was started.
<b>UpTime</b>	The amount of time, in milliseconds, since the server was started.
<b>maxMemory</b>	The maximum amount of memory used since the server was started.
<b>heapSize</b>	The heap size, in kilobytes.
<b>usedMemory</b>	The amount of used memory, in kilobytes.
<b>freeMemory</b>	The amount of free memory, in kilobytes.
<b>usedMemoryPercent</b>	The amount of used memory, in percent.
<b>HeapSize</b>	The heap size, in megabytes.
<b>UsedMemory</b>	The amount of used memory, in megabytes.
<b>FreeMemory</b>	The amount of free memory, in megabytes.

<b>maxHeapDumpsOnDisk</b>	The maximum amount of heap dumps on disk that have been performed since the last restart.
<b>maxSystemDumpsOnDisk</b>	The maximum amount of system dumps on disk that have been performed since the last restart.
<b>javaVendor</b>	The name of the Java software vendor.
<b>javaVersion</b>	The Java software version.
<b>TIME_STAMP</b>	The date and time of the last data update.

## All Servers Grid

View WebSphere server utilization details, including memory, CPU and heap size, in a grid format. Choose a cell and node from the drop-down menus.

Drill-down and investigate by clicking a row in the table to view details for the selected connection in the “[Server Summary](#)” display.



### Title Bar (possible features are):

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu ↓, Table open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

### Filter By:

- Cell:** Choose a cell or **All Cells** to see metrics for.
- Node:** Choose a node or **All Nodes** to see metrics for.

### Fields and Data

This display includes:

- Active Servers Only** Select to only include active servers in the display.
- Server Count** The number of servers in the display.
- Sort By:** Options are to sort servers in the grid by **Name**, **Live Sessions**, **JVM CPU %**, **Up Time** or **Max Memory**.

**Sort Descending** Select to organize display elements in descending order.

**Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

**Node** The name of the node.

**Cell** The name of the cell.

**UpTime** The amount of time, in milliseconds, since the server was started.

**PID** The WebSphere server process identifier.

**State** The WebSphere server current state:

- STARTED
- STOPPED

**Sessions** The current number of sessions.

**Phys Mem** The current amount of disk space, in megabytes.

**Trend Chart**

- **CPU** Traces the amount of server CPU utilization.
- **JVM** Traces the amount of server CPU utilization.
- **HEAP** Traces the amount of server heap memory utilization.
- **MEM** Traces the amount of server memory utilization.

## All Servers Detail

View detailed data for all your WebSphere servers in a tabular format, including system metrics, JVM runtime data and session totals.

Drill-down and investigate by clicking a row in the table to view details for a server in the “Server Summary” display.

← Grid All WebSphere Servers - Detail Tables 28-Nov-2016 10:39 Data OK + ?

**Server Info**

node	Server	Expired	state	cell	cellName	cellShortNam
TESTBED-27Node01	server1		STARTED	TESTBED-27Node01Cell	TESTBED-27Node01Cell	
TESTBED-28Node01	server1		STARTED	TESTBED-28Node01Cell	TESTBED-28Node01Cell	

System Metrics

node	Server	Expired	CPUUsageSinceLastMeasurement	CPUUsageSinceServerStarted
TESTBED-27Node01	server1		4	0 TESTBED
TESTBED-28Node01	server1		7	0 TESTBED

JVM Runtime Data

node	Server	Expired	UpTime	ProcessCpuUsage	HeapSize	UsedMemory	heapSize	n
TESTBED-27Node01	server1		8471851	2	112064	65667	114753536	
TESTBED-28Node01	server1		8471913	4	79907	57899	81824768	

Session Totals

node	Server	Expired	LiveCount	TIME_STAMP
TESTBED-27Node01	server1		0	11/28/16 1
TESTBED-28Node01	server1		0	11/28/16 1

#### Title Bar (possible features are):

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu ▾, Table open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

#### Server Info Table

Each table row is a different server. Table column values describe the server.

<b>Node</b>	The name of the node.
<b>Server</b>	The name of the server.
<b>Expired</b>	<ul style="list-style-type: none"> <li> Red indicates that the server is expired.</li> <li> Green indicates that the server is online.</li> </ul>
<b>State</b>	The WebSphere server current state: <ul style="list-style-type: none"> <li>• STARTED</li> <li>• STOPPED</li> </ul>
<b>cell</b>	The name of the cell.
<b>cellName</b>	The full name of the cell.

<b>cellShortName</b>	The short name for the cell.
<b>deployedObjects</b>	A list of deployed objects on the server.
<b>eventTypes</b>	A list of events that occurred on the server.
<b>internalClassAccessMode</b>	Describes the internal class access mode. Refer to vendor documentation for details.
<b>j2eeType</b>	The J2EE type.
<b>javaVMs</b>	A list of Java virtual machines.
<b>mbeanIdentifier</b>	The MBean id.
<b>name</b>	Refer to vendor documentation for details.
<b>nodeName</b>	The full name of the node.
<b>nodeShortName</b>	The short name for the node.
<b>pid</b>	The WebSphere server process identifier.
<b>platform</b>	The platform type.
<b>platformName</b>	The name of the platform.
<b>platformVersion</b>	The software version on the platform.
<b>processType</b>	The platform type (e.g. UnManagedProcess).
<b>resources</b>	Refer to vendor documentation for details.
<b>serverVendor</b>	The name of the server vendor.
<b>shortName</b>	Refer to vendor documentation for details.
<b>spec</b>	Refer to vendor documentation for details.
<b>statisticsProvider</b>	Refer to vendor documentation for details.
<b>threadMonitorAdjustmentThreshold</b>	Describes thread monitor settings. Refer to vendor documentation for details.
<b>threadMonitorInterval</b>	Describes thread monitor settings. Refer to vendor documentation for details.
<b>threadMonitorThreshold</b>	Describes thread monitor settings. Refer to vendor documentation for details.
<b>type</b>	Refer to vendor documentation for details.
<b>version</b>	Refer to vendor documentation for details.
<b>Connection</b>	The name of the connection.
<b>TIME_STAMP</b>	The date and time of the last data update.

**System Metrics Table**

Each table row is a different server. Table column values describe the server.

<b>Node</b>	The name of the node.
<b>Server</b>	The name of the server.
<b>Expired</b>	<ul style="list-style-type: none"> <li> Red indicates that the server is expired.</li> <li> Yellow indicates</li> <li> Green indicates that the server is online.</li> </ul>

<b>CPUUsageSinceLastMeasurement</b>	The amount of CPU usage, in megabytes, since the last data update.
<b>CPUUsageSinceServerStarted</b>	The amount of CPU usage, in megabytes, since the server was started.
<b>cell</b>	The name of the cell.
<b>hasStats</b>	Refer to vendor documentation for details.
<b>mbeanIdentifier</b>	The MBean id.
<b>name</b>	Refer to vendor documentation for details.
<b>platform</b>	Refer to vendor documentation for details.
<b>spec</b>	Refer to vendor documentation for details.
<b>type</b>	Refer to vendor documentation for details.
<b>Connection</b>	The name of the connection.
<b>TIME_STAMP</b>	The date and time of the last data update.

**JVM Runtime Data Table**

Each table row is a different server. Table column values describe the server.

<b>Node</b>	The name of the node.
<b>Server</b>	The name of the server.
<b>Expired</b>	 Red indicates that the server is expired.  Green indicates that the server is online.
<b>UpTime</b>	The amount of time, in milliseconds, since the server was started.
<b>ProcessCpuUsage</b>	The amount of process CPU usage, in megabytes, since the server was started.
<b>HeapSize</b>	The current heap size, in kilobytes.
<b>UsedMemory</b>	The current amount of memory used, in kilobytes.
<b>heapSize</b>	The current heap size, in kilobytes.
<b>maxMemory</b>	The maximum amount of memory used since the server was started.
<b>freeMemory</b>	The current amount of free memory, in kilobytes.
<b>usedMemory</b>	The current amount of used memory, in kilobytes.
<b>J2EEServer</b>	The name of the J2EE server.
<b>cell</b>	The name of the cell.
<b>hasStats</b>	Refer to vendor documentation for details.
<b>j2eeType</b>	The J2EE type.
<b>Java Vendor</b>	The name of the Java vendor.
<b>Java Version</b>	The Java software version.
<b>maxHeapDumpsOnDisk</b>	The maximum amount of heap dumps on disk that have been performed since the last restart.

<b>maxSystemDumpsOnDisk</b>	The maximum amount of system dumps on disk that have been performed since the last restart.
<b>mbeanIdentifier</b>	The MBean id.
<b>statisticsProvider</b>	Refer to vendor documentation for details.
<b>version</b>	Refer to vendor documentation for details.
<b>Connection</b>	The name of the connection.
<b>spec</b>	Refer to vendor documentation for details.
<b>platform</b>	Refer to vendor documentation for details.
<b>TIME_STAMP</b>	The date and time of the last data update.

**Session Totals Table**

Each table row is a different server. Table column values describe the server.

<b>Node</b>	The name of the node.
<b>Server</b>	The name of the server.
<b>Expired</b>	 Red indicates that the server is expired.  Yellow indicates  Green indicates that the server is online.
<b>LiveCount</b>	The total number of currently active sessions.
<b>TIME_STAMP</b>	The date and time of the last data update.

**Single Server View**

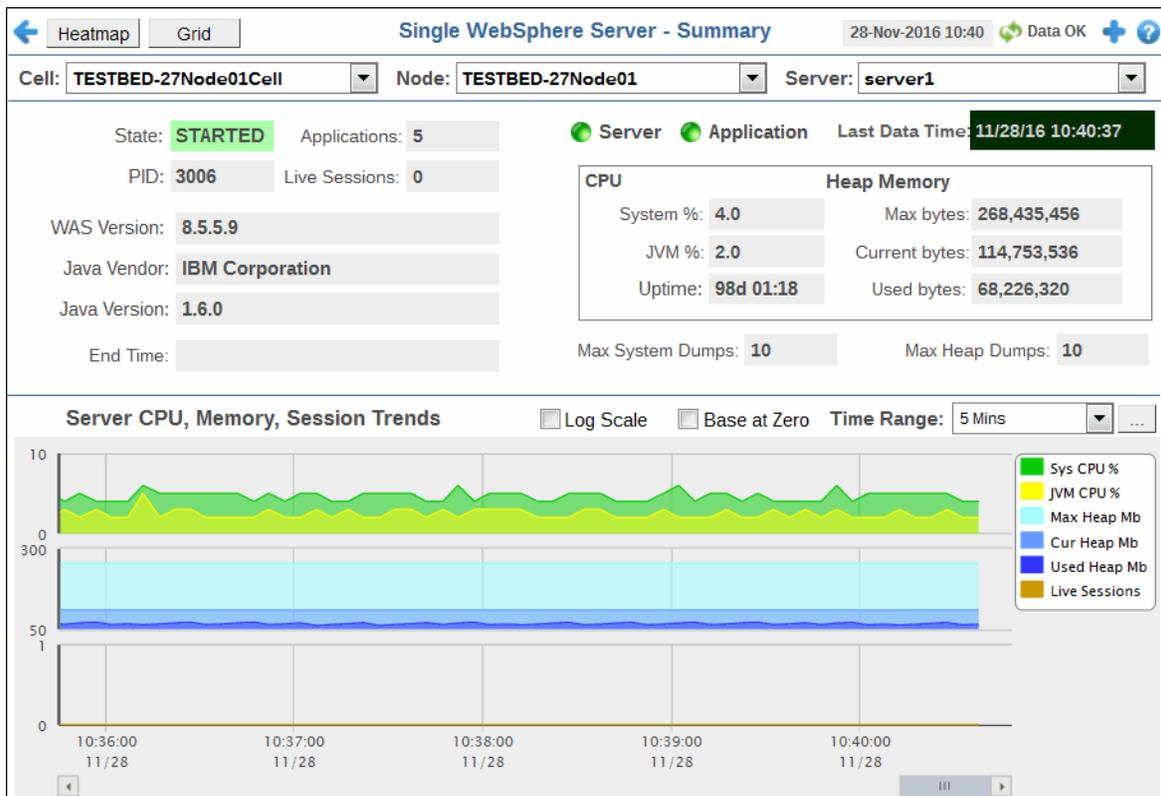
See performance and utilization metrics for a single IBM WebSphere server.

Displays in this View are:

- ["Server Summary"](#)
- ["JDBC Providers"](#)
- ["Thread Pool"](#)

## Server Summary

Track current and historical performance of web applications on one server. Choose a cell, node and server from the drop-down menus. Mouse over the trend graph to see metrics.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu Table open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

### Filter By:

- Cell:** Choose a cell or **All Cells** to see metrics for.
- Node:** Choose a node or **All Nodes** to see metrics for.
- Server:** Choose the server for which you want to show data.

### Fields and Data

This display includes:

- State** The WebSphere server current state:
  - STARTED
  - STOPPED

<b>Applications</b>	The number of applications running on the server.
<b>PID</b>	The WebSphere server process identifier.
<b>Live Sessions</b>	The current number of active sessions.
<b>WAS Version</b>	The WebSphere Application Server software version.
<b>Java Vendor</b>	The Java vendor name.
<b>Java Version</b>	The Java software version.
<b>End Time</b>	Refer to vendor documentation for details.
<b>Server</b>	The current alert severity.  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
<b>Application</b>	The current alert severity.  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
<b>Last Data Time</b>	The date and time of the last data update.

**CPU**

<b>System %</b>	The current amount of CPU used by the system, in percent.
<b>JVM %</b>	The current amount of CPU used by the JVM, in percent.
<b>Uptime</b>	The number of days, hours and minutes since the server was last started.

**Heap Memory**

<b>Max bytes</b>	The maximum size of heap memory, in bytes.
<b>Current bytes</b>	The current size of heap memory, in bytes.
<b>Used bytes</b>	The size of heap memory being used, in bytes.

**Max System Dumps** The maximum number of system dumps that are allowed to be performed.

**Max Heap Dumps** The maximum number of heap dumps that are allowed to be performed.

**Server CPU, Memory, Session Trends**

The trend graph traces the following for the selected server:

- **Sys CPU%** The percent of system CPU used.
- **JVM CPU%** The percent of JVM CPU used.
- **Max Heap Mb** The maximum amount of heap memory ever used, in megabytes.
- **Cur Heap Mb** The current amount of heap memory available, in megabytes.
- **Used Heap Mb** The current amount of heap memory used, in megabytes.
- **Live Sessions** The current number of active sessions.

**Log Scale**

Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.

**Base at Zero**

Use zero as the Y axis minimum for all graph traces.

**Time Range**

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .

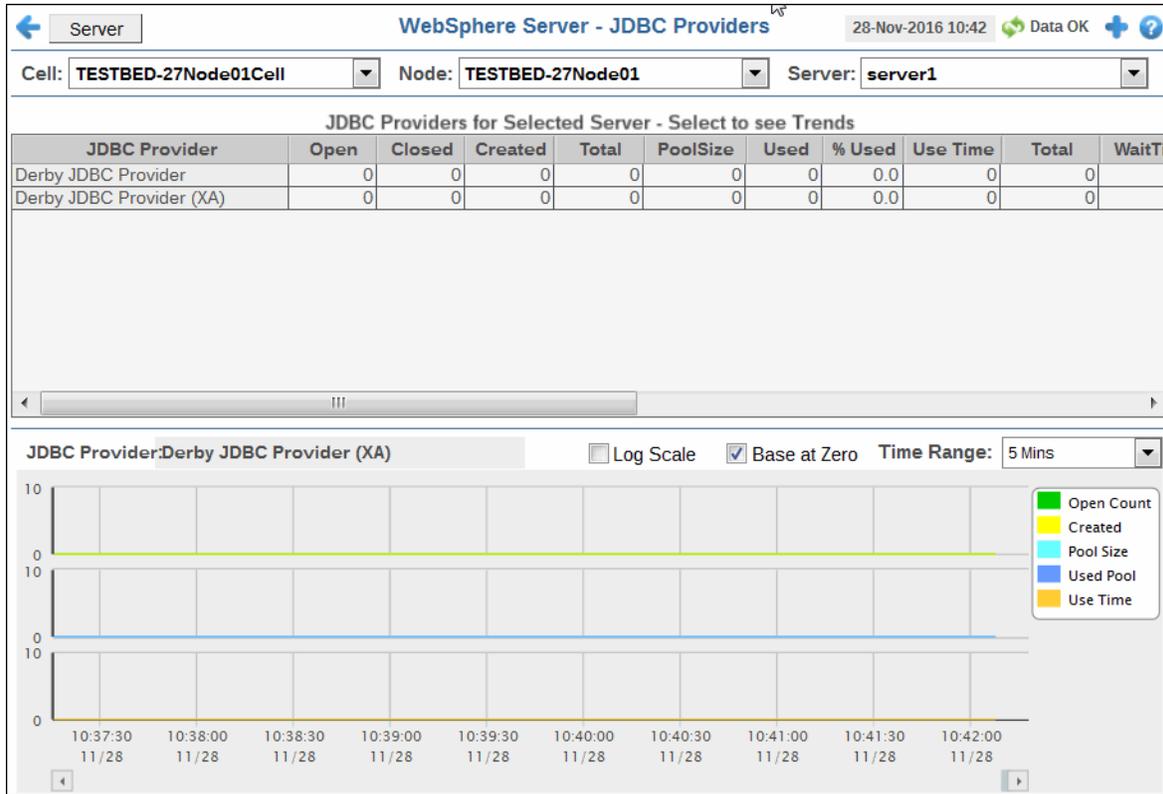
By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## JDBC Providers

Track current and historical performance of all JDBC Providers on a server. Choose a cell, node and server from the drop-down menus. Each table row is a different JDBC Provider. Select a row to populate the trend graph with JDBC Provider performance metrics. Mouse over the trend graph to see metrics.



**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** , **Table** open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

### Filter By:

- Cell:** Choose the cell for which you want to show data.
- Node:** Choose the node for which you want to show data.
- Server:** Choose the server for which you want to show data.

### JDBC Providers Table

Each table row is a different JDBC Provider.

- JDBC Provider** The name of the JDBC Provider.

<b>Open</b>	The number of currently open connections.
<b>Closed</b>	The number of currently closed connections.
<b>Created</b>	The number of connections created since the server was last started.
<b>Total</b>	Refer to vendor documentation for details.
<b>PoolSize</b>	The number of connections in the pool.
<b>Used</b>	The number of used connections in the pool.
<b>% Used</b>	The percent of connections used in the pool.
<b>Use Time</b>	The average connection duration, in seconds.
<b>Total</b>	Refer to vendor documentation for details.
<b>WaitTime</b>	The average amount of time to establish a connection, in seconds.
<b>WaitingThreadCount</b>	The current number of threads waiting to establish a connection, in seconds.
<b>description</b>	Describes the JDBC provider.
<b>TIME_STAMP</b>	The date and time of the last data update.

**JDBC Provider Trend Graph**

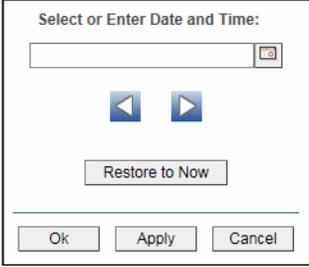
The trend graph traces the following for the JDBC Provider on the selected server:

- **Open Count** The number of currently open connections.
- **Created** The number of connections created.
- **Pool Size** The number of connections in the pool.
- **Used Pool** The number of connections in the pool that are being used.
- **Use Time** The average length of time of connections, in seconds.

**Log Scale** Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.

**Base at Zero** Use zero as the Y axis minimum for all graph traces.

**Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



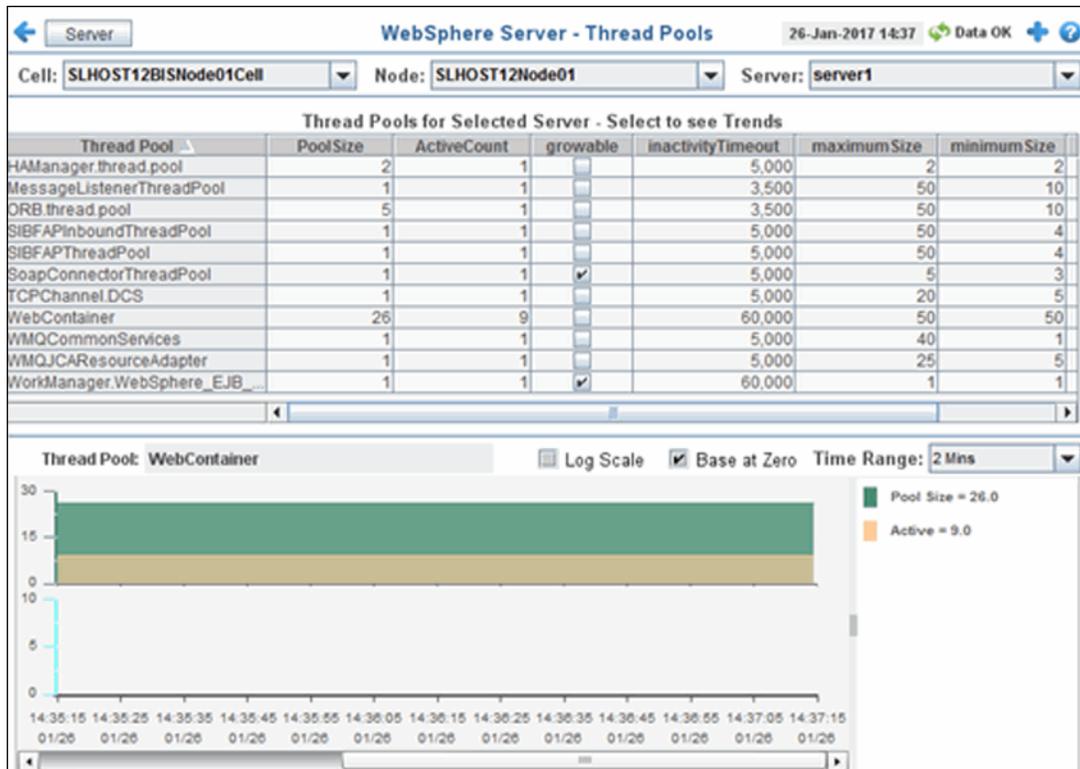
By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Thread Pool

Track current and historical performance of thread pools for web applications on one server. Choose a cell, node and server from the drop-down menus. Each table row is a different thread pool. Select a row to populate the trend graph with thread pool performance metrics. Mouse over the trend graph to see metrics.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

### Filter By:

- Cell:** Choose the cell for which you want to show data.
- Node:** Choose the node for which you want to show data.
- Server:** Choose the server for which you want to show data.

### Thread Pools Table

Each table row is a different thread pool.

- Thread Pool** The name of the thread pool.
- PoolSize** The number of threads in the pool.

<b>ActiveCount</b>	The number of currently active threads.
<b>growable</b>	When checked, the connection pool is growable. Refer to vendor documentation for details.
<b>InactivityTimeout</b>	Refer to vendor documentation for details.
<b>maximumSize</b>	Refer to vendor documentation for details.
<b>minimumSize</b>	Refer to vendor documentation for details.
<b>TIME_STAMP</b>	The date and time of the last data update.

### Thread Pool Trend Graph

The trend graph traces the following for the thread pool on the selected server:

- **Pool Size** The number of connections in the pool.
- **Active** The number of currently active connections in the pool.

**Log Scale** Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.

**Base at Zero** Use zero as the Y axis minimum for all graph traces.

**Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar  .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## All Applications View

See high-level performance and utilization metrics for all of your web application sessions.

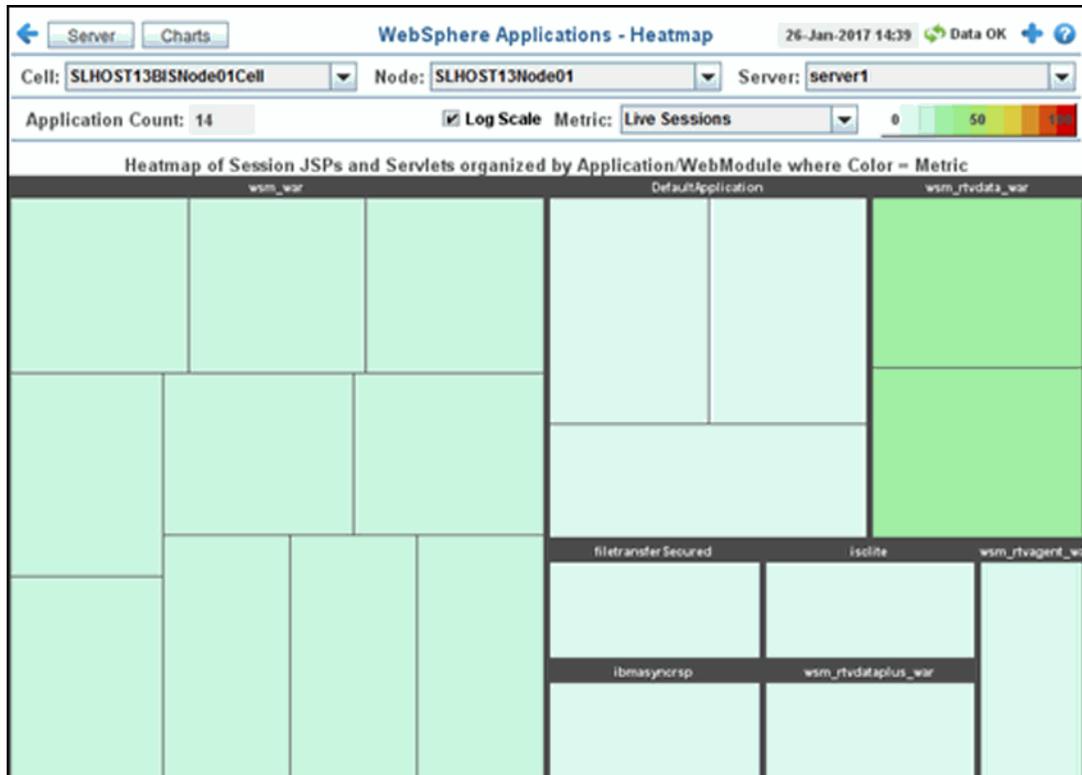
Displays in this View are:

- ["All Sessions Heatmap"](#)
- ["Session Charts By App"](#)
- ["Session Detail By App"](#)
- ["All Applications Detail"](#)

## All Sessions Heatmap

This heatmap shows activity metrics for all web application sessions on a selected server. Use this display to see metrics for **Live Sessions**, **Current Requests**, **Total Requests** and **Response Time** for all web application sessions on a server. By default, this display shows the heatmap based on the **Live Sessions** metric.

Each rectangle is a different web application on the server. Choose a **Cell**, **Node** and **Server** from the drop-down menus. Mouse over a rectangle to see additional metrics. Click a rectangle to drill down to the "Application Summary" display, which shows additional details about the application.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

### Filter By:

- Cell:** Choose the cell for which you want to show data.
- Node:** Choose the node for which you want to show data.
- Server:** Choose the server for which you want to show data.

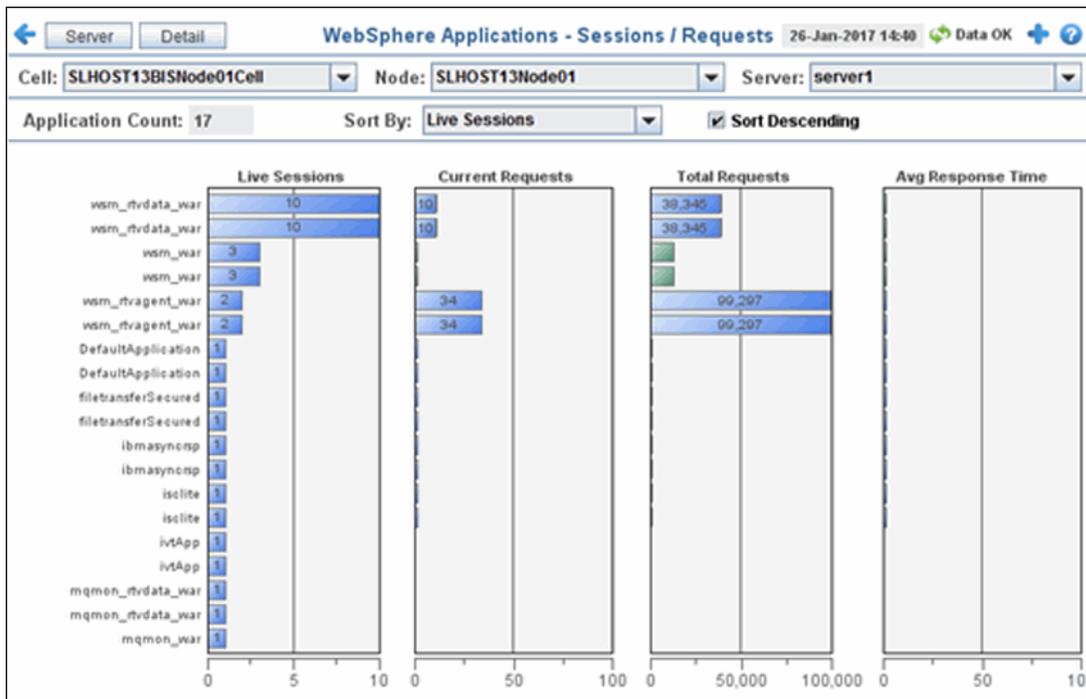
### Fields and Data:

<b>Application Count</b>	The number of web applications in the display.
<b>Log Scale</b>	Select to enable a logarithmic scale. Use <b>Log Scale</b> to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. <b>Log Scale</b> makes data on both scales visible by applying logarithmic values rather than actual values to the data.
<b>Metric</b>	Choose a metric to view in the display.
<b>Live Sessions</b>	The current number of live sessions. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of <b>WasLiveSessionCountHigh</b> . The middle value in the gradient bar indicates the middle value of the range.
<b>Current Requests</b>	The number of current requests. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of <b>WasServletRequestRateHigh</b> and <b>WasJspRequestRateHigh</b> . The middle value in the gradient bar indicates the middle value of the range.
<b>Total Requests</b>	The total number of requests. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum number of requests. The middle value in the gradient bar indicates the average amount.
<b>Response Time</b>	The average response time. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of <b>WasServletResponseTimeHigh</b> and <b>WasJspResponseTimeHigh</b> . The middle value in the gradient bar indicates the middle value of the range.

## Session Charts By App

This display shows activity metrics for all web application sessions on a selected server. Choose a cell, node and server from the drop-down menus. Use this display to see metrics for **Live Sessions, Current Servlet Requests, Total Servlet Requests and Total JSP Requests, Current JSP Requests** and average response times for all web application sessions on a server. By default, this display shows the **Live Sessions** metric.

Click on an object to drill down to details.



**Title Bar (possible features are):**

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu Table open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

**Filter By:**

- Cell:** Choose the cell for which you want to show data.
- Node:** Choose the node for which you want to show data.
- Server:** Choose the server for which you want to show data.

**Fields and Data:**

- Application Count** The number of web applications in the display.
- Sort By:** Choose a value to sort data in the display:
  - Name** The name of the item.
  - Live Sessions** The current number of active sessions.
  - Total Servlet Requests** The number of servlet requests since the server was started.

**Current Servlet Requests**

The current number of servlet requests.

**Total JSP Requests**

The number of JSP requests since the server was started.

**Current JSP Requests**

The current number of JSP requests.

**Servlet Avg Response Time**

The average amount of time for the servlet to respond.

**JSP Avg Response Time**

The average amount of time for the JSP to respond.

**Sort Descending**

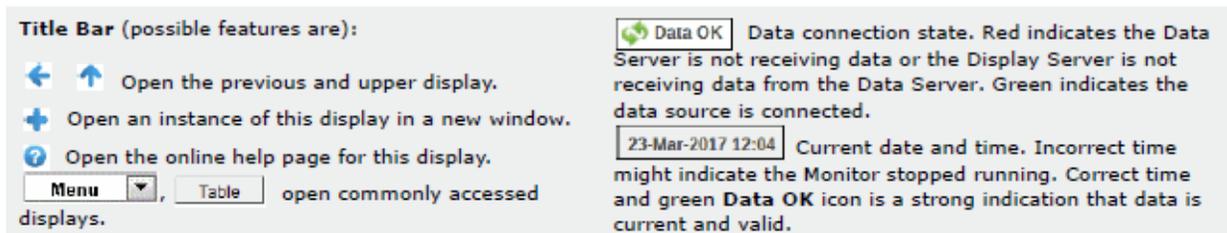
Select to sort data in descending order.

**Session Detail By App**

This display shows activity metrics for all web applications on a selected server. Choose a cell, node and server from the drop-down menus. Each row in the table is a different application.

The screenshot shows the 'WebSphere Applications - Sessions / Requests' interface. At the top, there are navigation buttons for 'Server' and 'Charts', and a title bar with the date '26-Jan-2017 14:41' and 'Data OK'. Below the title bar, there are three dropdown menus: 'Cell: SLHOST13BISNode01Cell', 'Node: SLHOST13Node01', and 'Server: server1'. A status bar indicates 'Application Count: 17'. The main content is a table titled 'All Applications for Selected Server' with the following columns: Application, Sessions, Servlets, Total Reqs, Current Reqs, Avg Resp Time, JSPs, Total Reqs, and Current Req. The table lists various applications such as wsm\_rtdata\_war, wsm\_war, wsm\_rtagent\_war, DefaultApplication, filetransferSecured, ibmasyncrsp, isclite, MApp, mqmon\_rtdata\_war, and SamplesGallery, along with their respective metrics.

Application	Sessions	Servlets	Total Reqs	Current Reqs	Avg Resp Time	JSPs	Total Reqs	Current Req
wsm_rtdata_war	10	1	38,345	10	0.2	1	870	
wsm_rtdata_war	10	1	38,345	10	0.2	1	870	
wsm_war	3	0	0	0	0.0	10	12,431	
wsm_war	3	0	0	0	0.0	10	12,431	
wsm_rtagent_war	2	1	99,297	34	1.1	0	0	
wsm_rtagent_war	2	1	99,297	34	1.1	0	0	
DefaultApplication	1	3	1	1	1.1	0	0	
DefaultApplication	1	3	1	1	1.1	0	0	
filetransferSecured	1	1	1	1	1.1	0	0	
filetransferSecured	1	1	1	1	1.1	0	0	
ibmasyncrsp	1	1	1	1	1.1	0	0	
ibmasyncrsp	1	1	1	1	1.1	0	0	
isclite	1	1	1	1	1.1	0	0	
isclite	1	1	1	1	1.1	0	0	
MApp	1	0	0	0	0.0	0	0	
MApp	1	0	0	0	0.0	0	0	
mqmon_rtdata_war	1	0	0	0	0.0	0	0	
mqmon_rtdata_war	1	0	0	0	0.0	0	0	
mqmon_war	1	0	0	0	0.0	0	0	
mqmon_war	1	0	0	0	0.0	0	0	
PlantsByWebSphere	1	0	0	0	0.0	0	0	
PlantsByWebSphere	1	0	0	0	0.0	0	0	
SamplesGallery	1	0	0	0	0.0	0	0	
SamplesGallery	1	0	0	0	0.0	0	0	
wsm_rtdataplus_war	1	1	1	1	1.1	0	0	
wsm_rtdataplus_war	1	1	1	1	1.1	0	0	

**Filter By:**

- Cell:** Choose the cell for which you want to show data.
- Node:** Choose the node for which you want to show data.
- Server:** Choose the server for which you want to show data.

**Fields and Data:**

- Application Count** The number of web applications in the display.

**All Applications Table**

Column values describe the application.

- Application** The name of the application.
- Sessions** The number of current sessions.
- Servlets** The number of servlets.
- Total Reqs** The number of requests since the application was started.
- Current Reqs** The number of current requests.
- Avg Resp Time** The average response time, in seconds.
- JSPs** The number of JSPs.
- Total Reqs** The number of requests since the application was started.
- Current Reqs** The number of current requests.
- Total Reqs** The number of requests since the application was started.
- Current Reqs** The number of current requests.
- Avg Resp Time** The average response time, in seconds.

## All Applications Detail

This display shows detailed application information and deployment descriptors for all web application sessions on a selected server. Choose a cell, node and server from the drop-down menus. Select an application in the upper table to see the deployment descriptor in the lower portion of the display.

**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

### Filter By:

- Cell:** Choose the cell for which you want to show data.
- Node:** Choose the node for which you want to show data.
- Server:** Choose the server for which you want to show data.

### Fields and Data:

- Application Count** The number of web applications in the display.

### Application Info Table

Each table row is a different application. Click a row to see data in the **Deployment Descriptor** table.

<b>Application</b>	The name of the application.
<b>implementationVersion</b>	The application version.
<b>statisticsProvider</b>	The name of the statistics provider.
<b>eventTypes</b>	A list of application events.
<b>version</b>	The application version.
<b>modules</b>	A list of application modules.
<b>TIME_STAMP</b>	The date and time of the last data update.

**Session Manager Data Table**

Each table row is a different application.

<b>Application</b>	The name of the application.
<b>LiveCount</b>	The number of connections for the application.
<b>TIME_STAMP</b>	The date and time of the last data update.

**Deployment Descriptor:**

Provides details about the application deployment.

## Single Application View

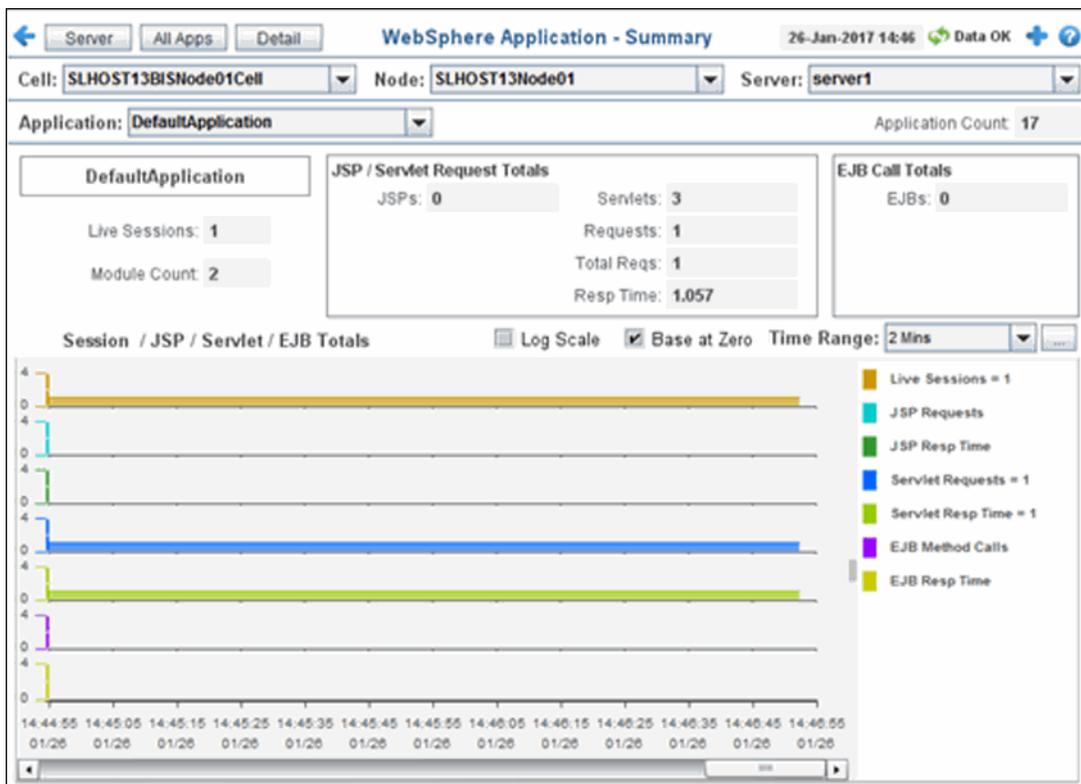
See performance and utilization metrics for a single web application.

Displays in this View are:

- ["Application Summary"](#)
- ["Component Detail"](#)
- ["Module Totals - Charts"](#)
- ["Module Totals - Tables"](#)

## Application Summary

View performance metrics for a web application on one server. Choose a cell, node, server and application from the drop-down menus.



**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Filter By:**

- Cell:** Choose the cell for which you want to show data.
- Node:** Choose the node for which you want to show data.
- Server:** Choose the server for which you want to show data.
- Application:** Choose the application for which you want to show data.

**Fields and Data**

This display includes:

- Application Count** The number of applications in the display.

**Live Sessions** The current number of active sessions.

**Module Count** The current number of modules.

### JSP/Servlet Request Totals

**JSPs** The current number of JSP requests.

**Servlets** The current number of servlet requests.

### EJB Call Totals

**EJBs** The current number of EJB (Enterprise JavaBean) requests.

### Session/JSP/Servlet/EJB Totals Trend Graph

The trend graph traces the following for the selected application:

- **Live Sessions** The current number of active sessions.
- **JSP Requests** The current number of JSP requests.
- **JSP Resp Time** The current average JSP response time, in seconds.
- **Servlet Requests** The current number of servlet requests.
- **Servlet Resp Time** The current average servlet response time, in seconds.
- **EJB Method Calls** The current number of EJB requests.
- **EJB Resp Time** The current average EJB response time, in seconds.

**Log Scale** Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.

**Base at Zero** Use zero as the Y axis minimum for all graph traces.

**Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .

By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Component Detail

View performance metrics for a web application on one server. Choose a cell, node, server and application from the drop-down menus.

The screenshot shows the 'WebSphere Application - Component Tables' interface. At the top, there are navigation buttons for 'All Apps' and 'Summary', and a title bar with the date and time '28-Nov-2016 13:29' and a 'Data OK' status. Below the title bar, there are several filter menus: 'Cell: TESTBED-27Node01Cell', 'Node: TESTBED-27Node01', 'Server: server1', and 'Application: DefaultApplication'. The 'Application Count' is shown as 8.

The main content area is divided into four sections:

- Live Sessions:** A table with columns: WebModule, LiveCount, NoRoomForNewSessionCour, LifeTime, ExternalRe, ExternalR, ExternalWr, ExternalWr. The row for 'DefaultWebApplication.war' shows all values as 0.
- JSPs:** A section with a grey header and an empty table below it.
- Servlets:** A table with columns: Name, Requests, Total Requests, Avg Response Time Rec, Total Response Time, Avg Response Tin. The rows are 'Hello Pervasive Servlet', 'Hit Count Servlet', and 'Snoop Servlet', all with 0 requests and 0.0 response times.
- EJBs:** A section with a grey header and an empty table below it.

**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

### Filter By:

- Cell:** Choose the cell for which you want to show data.
- Node:** Choose the node for which you want to show data.
- Server:** Choose the server for which you want to show data.
- Application:** Choose the application for which you want to show data.

### Fields and Data:

- Application Count** The number of applications in the display.

**Live Sessions Table**

Each table row is a different web module. Table columns describe the module.

<b>WebModule</b>	The name of the web module.
<b>LiveCount</b>	The current number of sessions.
<b>NoRoomForNewSessionCount</b>	The number of times the module had no sessions available.
<b>Life Time</b>	Refer to vendor documentation for details.
<b>ExternalReadSize</b>	Refer to vendor documentation for details.
<b>ExternalWriteSize</b>	Refer to vendor documentation for details.
<b>ExternalWriteTime</b>	Refer to vendor documentation for details.

**JSPs Table**

Each table row is a different JSP. Table columns describe the JSP.  
Refer to vendor documentation for details.

**Servlets Table**

Each table row is a different servlet. Table columns describe the servlet.

<b>Name</b>	The name of the servlet.
<b>Requests</b>	The current number of requests.
<b>Total Requests</b>	The total number of requests since the servlet was started.
<b>Avg Response Time Recent</b>	The current average amount of time for the servlet to respond.
<b>Total Response Time</b>	The total response time, in seconds, since the servlet was started.
<b>Avg Response Time</b>	The average amount of time for the servlet to respond since the servlet was started.

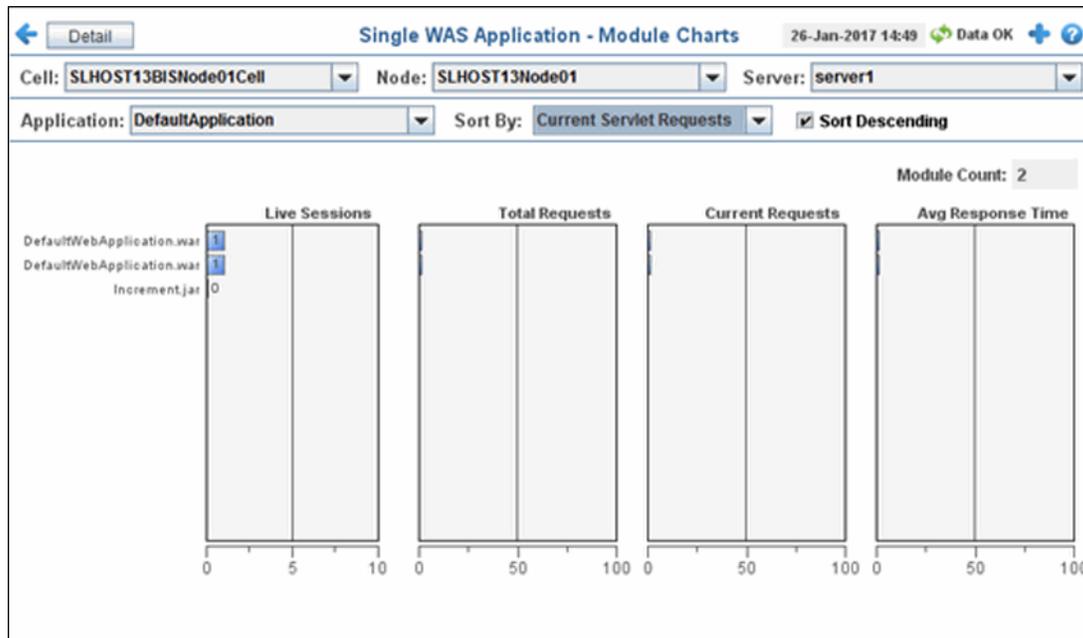
**EJBs Table**

Each table row is a different EJB. Table columns describe the EJB.

<b>Name</b>	The name of the EJB.
<b>CreateCount</b>	The number of requests.
<b>MethodCalls</b>	The total number of requests since the servlet was started.
<b>Total Calls</b>	The current average amount of time for the servlet to respond.
<b>Ready Count</b>	Refer to vendor documentation for details.
<b>Response Time</b>	The average amount of time for the EJB to respond.
<b>PassiveCount</b>	Refer to vendor documentation for details.
<b>RemoveCount</b>	Refer to vendor documentation for details.

## Module Totals - Charts

View performance metrics for a WAS application on one server. Choose a cell, node, server and application from the drop-down menus.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

### Filter By:

- Cell:** Choose the cell for which you want to show data.
- Node:** Choose the node for which you want to show data.
- Server:** Choose the server for which you want to show data.
- Application:** Choose the application for which you want to show data.
- Sort By:**
  - **Name**
  - **Live Sessions**
  - **Total Servlet Requests**
  - **Current Servlet Requests**
  - **Total JSP Requests**
  - **Current JSP Requests**
  - **Servlet Avg Resp Time**
  - **JSP Avg Resp Time**

**Fields and Data:**

**Sort Descending** Select to organize display elements in descending order.

**Module Count** The number of modules in the display.

**Graphs**

Refer to vendor documentation for details.

**Live Sessions Graph** Shows performance metrics for current sessions.

**Total Requests Graph** Shows performance metrics for total requests.

**Current Requests Graph** Shows performance metrics for current requests.

**Avg Response Time Graph** Shows performance metrics for average response time.

**Module Totals - Tables**

View performance metrics for a web application on one server. Choose a cell, node, server and application from the drop-down menus. Each row in the upper table is a different module for the selected application. Select a row to populate the lower tables.

**Single WAS Application - Module Totals** 26-Jan-2017 14:51 Data OK

Cell: SLHOST13Node01Cell Node: SLHOST13Node01 Server: server1

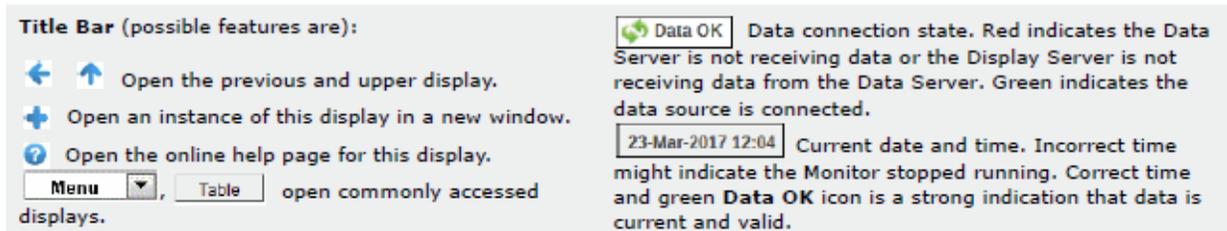
Application: iscite **Module Count: 6**

Module Name	Context Root	Sessions	Servlets	Total Reqs	Current Reqs	Avg Resp Time	JSPs
adminredirector.war	admin_host/admin	1	0	0	0	0.0	0
adminredirector.war	admin_host/admin	1	0	0	0	0.0	0
jehs.war	admin_host/ibm/help	1	0	0	0	0.0	0
jehs.war	admin_host/ibm/help	1	0	0	0	0.0	0
ISCAdminPortlet.war	admin_host/ISCAdminPortlet	1	0	0	0	0.0	0
ISCAdminPortlet.war	admin_host/ISCAdminPortlet	1	0	0	0	0.0	0
iscite.war	admin_host/ibm/console	1	1	1	1	1.1	0
iscite.war	admin_host/ibm/console	1	1	1	1	1.1	0
wasportlet.war	admin_host/wasportlet	1	0	0	0	0.0	0
wasportlet.war	admin_host/wasportlet	1	0	0	0	0.0	0
WIMPortlet.war	admin_host/wim	1	0	0	0	0.0	0
WIMPortlet.war	admin_host/wim	1	0	0	0	0.0	0

**EJBs for Selected Module**

**Portlets for Selected Module**

- iscite#ISCAdminPortlet.war\_portlet.ConsoleIdentityPortlet
- iscite#ISCAdminPortlet.war\_portlet.ISCProductDetails
- iscite#ISCAdminPortlet.war\_portlet.PortletPref
- iscite#ISCAdminPortlet.war\_portlet.WelcomePortlet

**Filter By:**

- Cell:** Choose the cell for which you want to show data.
- Node:** Choose the node for which you want to show data.
- Server:** Choose the server for which you want to show data.
- Application:** Choose the application for which you want to show data.

**Fields and Data:**

**Module Count** The number of modules in the display.

**All Modules Table**

Each table row is a different module. Column values describe the module. Select a module to see EJBs and Portlets for the module in the lower tables.

<b>Module Name</b>	The name of the module.
<b>Context Root</b>	<b>The context root.</b>
<b>Sessions</b>	The current number of sessions.
<b>Servlets</b>	The current number of servlets.
<b>Total Requests</b>	The total number of requests since the servlet was started.
<b>Current Requests</b>	The current number of requests.
<b>Avg Response Time</b>	The average amount of time to respond, in seconds.
<b>JSPs</b>	The current number of JSPs.
<b>Total Response Time</b>	The total response time, in seconds, since the servlet was started.
<b>Avg Response Time</b>	The average amount of time for the servlet to respond since the servlet was started.
<b>Total Requests</b>	The total number of requests since the servlet was started.
<b>Current Requests</b>	The current number of requests.

**Avg  
Response  
Time**

The average amount of time to respond, in seconds.

**EJBs for Selected Module**

List of EJBs for the selected module.

**Portlets for Selected Module**

List of portlets for the selected module.

---

## IBM WebSphere - HTML

The IBM WebSphere HTML displays provide extensive visibility into the health and performance of IBM WebSphere application servers and installed web modules. The following IBM WebSphere Views (and their associated displays) can be found under **Components** tab > **Application/Web Servers** > **IBM WebSphere**.

IBM WebSphere has the following displays:

- ["WebSphere Overview - HTML"](#)
- ["WebSphere Servers Heatmap - HTML"](#): Performance metrics for one IBM WebSphere Server, including current and historic performance metrics.
- ["WebSphere Server Summary - HTML"](#): Heatmap of performance metrics for all Web modules for one IBM WebSphere Server.
- ["WebSphere Apps Table - HTML"](#): Table and trend graphs of performance metrics for Web modules.
- ["WebSphere Apps Heatmap - HTML"](#)
- ["WebSphere App Summary - HTML"](#)

### WebSphere Overview - HTML

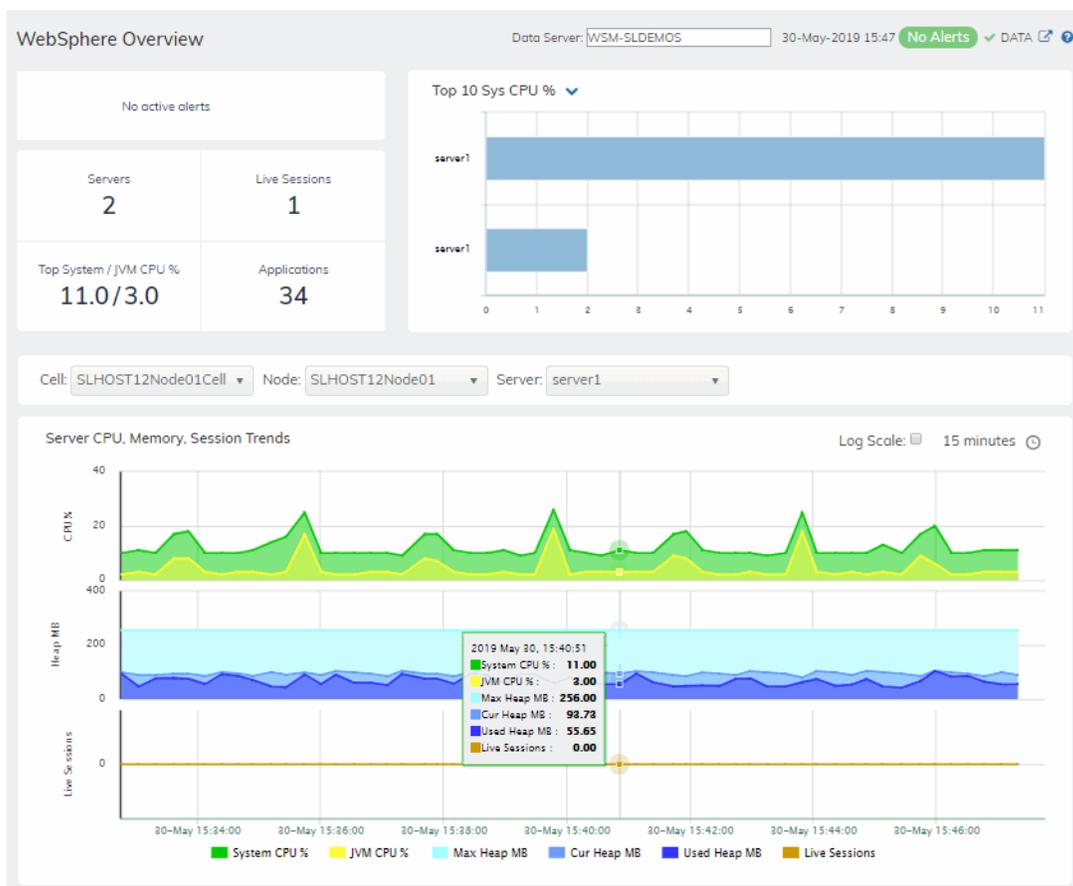
The WebSphere Overview is the top-level display for the WebSphere Solution Package, which provides a good starting point for immediately getting the status of all your WebSphere servers, web modules and connections. You can select the RTView DataServer for which you want to see data and easily view the current data for that DataServer including:

- The total number of active alerts for the selected DataServer, including the total number of critical and warning alerts.
- The greatest number of **Live Sessions** and **Top System/JVM CPU%**.
- A bar graph shows the WebSphere servers with the **Top 10 System CPU %** usage, and allows you to instead show the **Top 10 JVM CPU %** usage, the **Top 10 System Max Heap MB** usage, the **Top 10 Cur Heap** usage, the **Top 10 Used Heap** usage or the **Top 10 Live Sessions**.

You can hover over the metric cards in the upper half of the Overview and click to investigate details in a Summary display.

You can choose a **Cell**, **Node** and **Server** from the drop-down menus for the trend graph which traces **System CPU %**, **Live Sessions**, **JVM CPU%**, **Max Heap** and **Used Heap MB**.

You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.



## WebSphere Servers Table - HTML

Investigate detailed utilization metrics and configuration settings for all or one WebSphere cell and all or one WebSphere node. The **WebSphere Servers Table** contains all metrics available for servers, including the number of current client connections.

Each row in the table contains data for a particular **Node**. You can search, filter, sort and choose columns to include by clicking a column header icon (to the right of each column label) and selecting **Filter**, **Sort Ascending**, **Sort Descending** or **Columns**. Or just click a column header to sort.

Right-click on a table cell to **Export to Excel** or **Copy Cell Value**.

You can click on a row to drill-down to the ["WebSphere Server Summary - HTML"](#) display and view details for that server.

## WebSphere Servers Heatmap - HTML

View performance metrics for all monitored WebSphere Servers. The heatmap organizes WebSphere Web modules by server, and uses color to show the most critical Metric value for each WebSphere connection associated with the selected source. Each rectangle in the heatmap represents a Web module. In this heatmap, the rectangle size represents the value for maximum heap memory used. Each **Metric** (selected from the drop-down menu) has a color gradient bar that maps relative values to colors.

Use this display to see at-a-glance the health of all your web applications. You can select the heatmap color metric from a list including active sessions, access rate, and total access count.

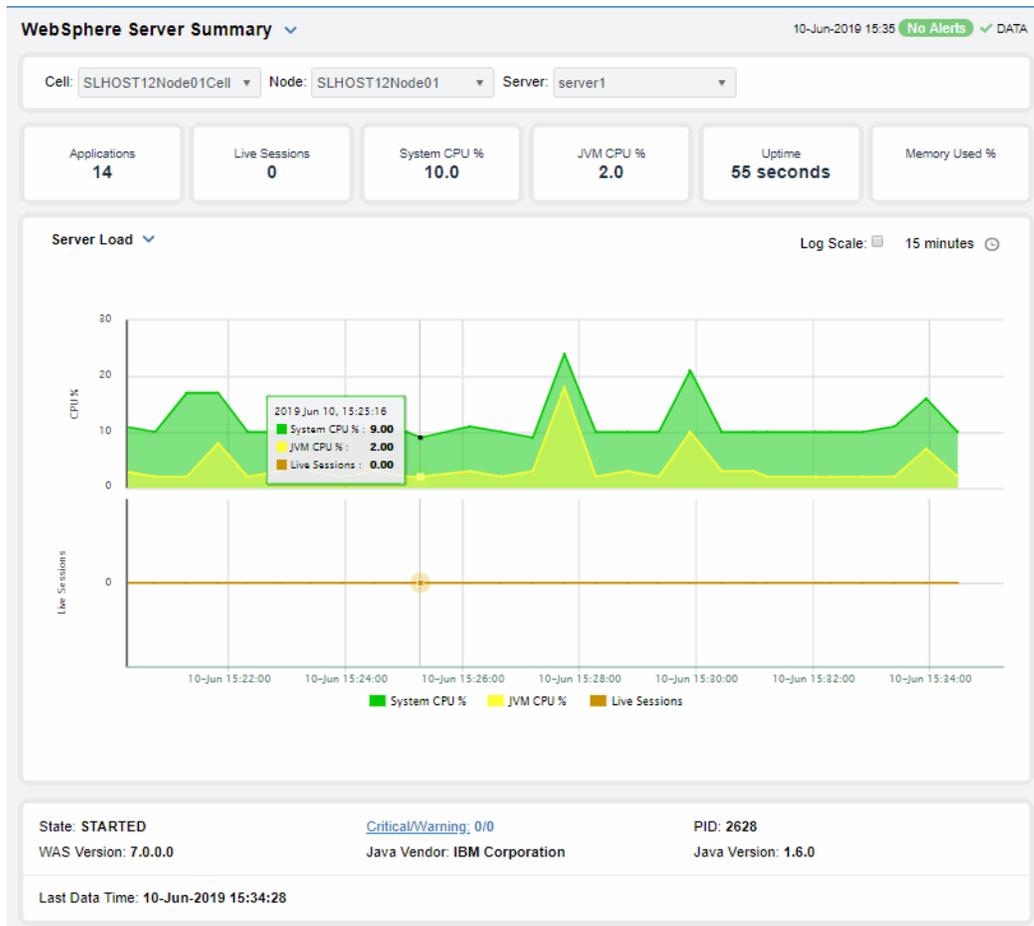
Use the available drop-down menus or right-click to filter data shown in the display. Use the check-boxes  to include or exclude labels in the heatmap. Move your mouse over a rectangle to see additional information. Drill-down and investigate by clicking a rectangle in the heatmap to view details for the selected Web module in the **Server Summary** display.

## WebSphere Server Summary - HTML

Track utilization and performance metrics for a connection on a WebSphere server. Clicking on the sessions/processing rate information boxes at the top of the display takes you to the **WebSphere Servers Table** display, where you can compare and sort performance values against other WebSphere servers.

The trend graph traces for **Processing Time per second**, **Requests per second** and (number of) **Active Sessions**. You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.

Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display.



## WebSphere Apps Table - HTML

Investigate detailed utilization metrics for all WebSphere applications. This display contains all metrics available for WebSphere applications, including the total **Alert Count**, **Accesses/per second** and **Total Sessions**.

Choose a particular **Source** or **All**, and a particular **Connection** or **All**, from the drop-downs. Each row in the table contains data for a particular web module. You can search, filter, sort and choose columns to include by clicking a column header icon (to the right of each column label) and selecting **Filter**, **Sort Ascending**, **Sort Descending** or **Columns**.

To investigate further, double-click a web module to see details in the **WebSphere Application Summary display**.

Or just click a column header to sort. Right-click on a table cell to **Export to Excel** or **Copy Cell Value**.

## WebSphere Apps Heatmap - HTML

This heatmap allows you to view the status and alerts of WebSphere applications on a particular host or **All** hosts, and a particular connection or **All** connections.

Use the **Metric** drop-down menu to view the **Alert Severity, Alert Count, Active Sessions, Accesses per Second** or (the total number of) **Accesses**.

Each rectangle in the heatmap represents a web module. The rectangle color indicates the most critical alert state. Click on a rectangle to drill-down to the **WebSphere Application Summary** display and view metrics for a particular web module. Toggle between the commonly accessed Table and Heatmap displays by clicking the drop down list on the display title.

Mouse-over rectangles to view more details about host performance and status.

You can view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.

## WebSphere App Summary - HTML

Track utilization and performance metrics for a particular WebSphere web module. Clicking on the sessions/processing rate information boxes at the top of the display takes you to the **WebSphere Servers Table** display, where you can compare and sort performance values against other WebSphere servers.

Use the **Web Modules** table to compare detailed utilization metrics for all web modules. Each row in the table contains data for a particular web module. You can search, filter, sort and choose columns to include by clicking a column header icon (to the right of each column label) and selecting **Filter, Sort Ascending, Sort Descending** or **Columns**.

Or just click a column header to sort. Right-click on a table cell to **Export to Excel** or **Copy Cell Value**.

The trend graph traces for **Processing Time per second, Accesses per second** and (the number of) **Active Sessions**. You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.

Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display.

## CHAPTER 4 RTView DataServer for Infrastructure

The RTView DataServer for Infrastructure provides a way to create connections and modify default configuration settings for solution packages and sends collected data to RTViewCentral. RTViewCentral contains the displays associated with the RTView DataServer for Infrastructure which you use to monitor your Infrastructure components. Both the display server user interface and the HTML user interface are described here.

For an overview and details about configuring RTView Enterprise, including RTViewCentral, RTView DataServers, RTView DataCollectors and solution packages, see the *RTView Enterprise Configuration Guide*.

RTViewCentral contains the following solution packages and associated displays that will be populated with data collected via the RTView DataServer for Infrastructure:

- ["Amazon Web Services"](#): Describes the Display Server version.
- ["Amazon Web Services - HTML"](#): Describes the HTML version.
- ["Docker"](#): Describes the display server version.
- ["Docker - HTML"](#): Describes the HTML version.
- ["JBoss"](#): Describes the display server version.
- ["JBoss - HTML"](#): Describes the HTML version.
- ["MongoDB"](#): Describes the display server version.
- ["MongoDB - HTML"](#): Describes the HTML version.
- ["MySQL Database"](#): Describes the display server version.
- ["MySQL Database - HTML"](#): Describes the HTML version.
- ["MS SQL Server"](#): Describes the display server version.
- ["MS SQL - HTML"](#): Describes the HTML version.
- ["Node.js"](#): Describes the display server version.
- ["Node.js - HTML"](#): Describes the HTML version.
- ["RTView Host Agent"](#): Describes the display server version.
- ["RTView Host Agent - HTML"](#): Describes the HTML version.
- ["VMware vCenter"](#): Describes the display server version.
- ["VMware vCenter - HTML"](#): Describes the display server version.

The following displays are also impacted by the settings in the RTView DataServer for Infra:

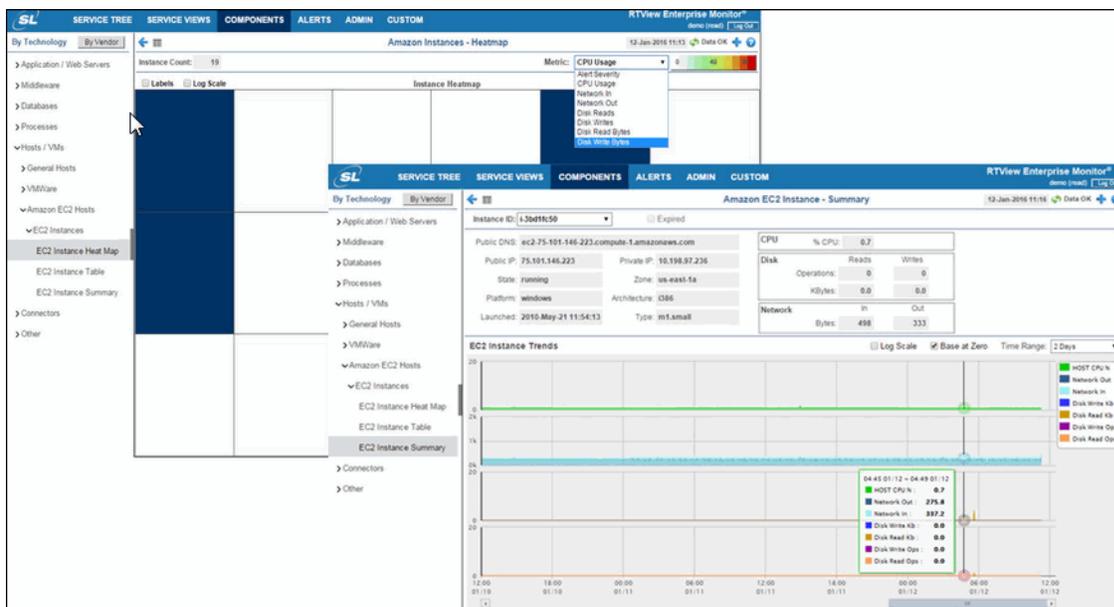
- ["JVM Displays"](#)
- ["Tomcat Displays"](#)

The RTView *DataCollector* for Infra is also available for use with the RTView DataServer for IBM. RTView DataCollector for Infra is used for collecting and sending data to one or more data servers. The RTView DataCollector for Infra is also useful if you need to distribute data collection.

**Note:** This document assumes familiarity with the products monitored. For additional details, refer to vendor documentation.

## Amazon Web Services

The Solution Package for Amazon Web Services provides a high level Amazon Instance Heatmap for a complete view of your AWS infrastructure with drill down views to individual AWS instances.



Using the RTView Historian, Amazon AWS metrics are persisted to a database for trend analysis. Historical trends are then used to help define alert thresholds against Amazon AWS data which, when correlated with alerts from other application components through RTView Enterprise's alert management system, can help users identify the source of performance problems more quickly.

With the Solution Package for Amazon Web Services, you are able to drill down from a high level alert on a business service or application into the supporting Amazon AWS infrastructure to determine what is causing the alert and take corrective action. This service-centric approach makes it easy for application support teams to prioritize incidents based on the impact to the business.

The following Solution Package for Amazon Web Services Views (and associated displays) can be found under **Components tab > Hosts/VMs > Amazon EC2 Hosts**. For additional details, see vendor documentation.

This section contains the following:

- "EC2 Instances"

## EC2 Instances

Displays in this View are:

- ["Amazon EC2 Instance Heatmap"](#)
- ["Amazon EC2 Instance Table"](#)
- ["Amazon EC2 Instance Summary"](#)

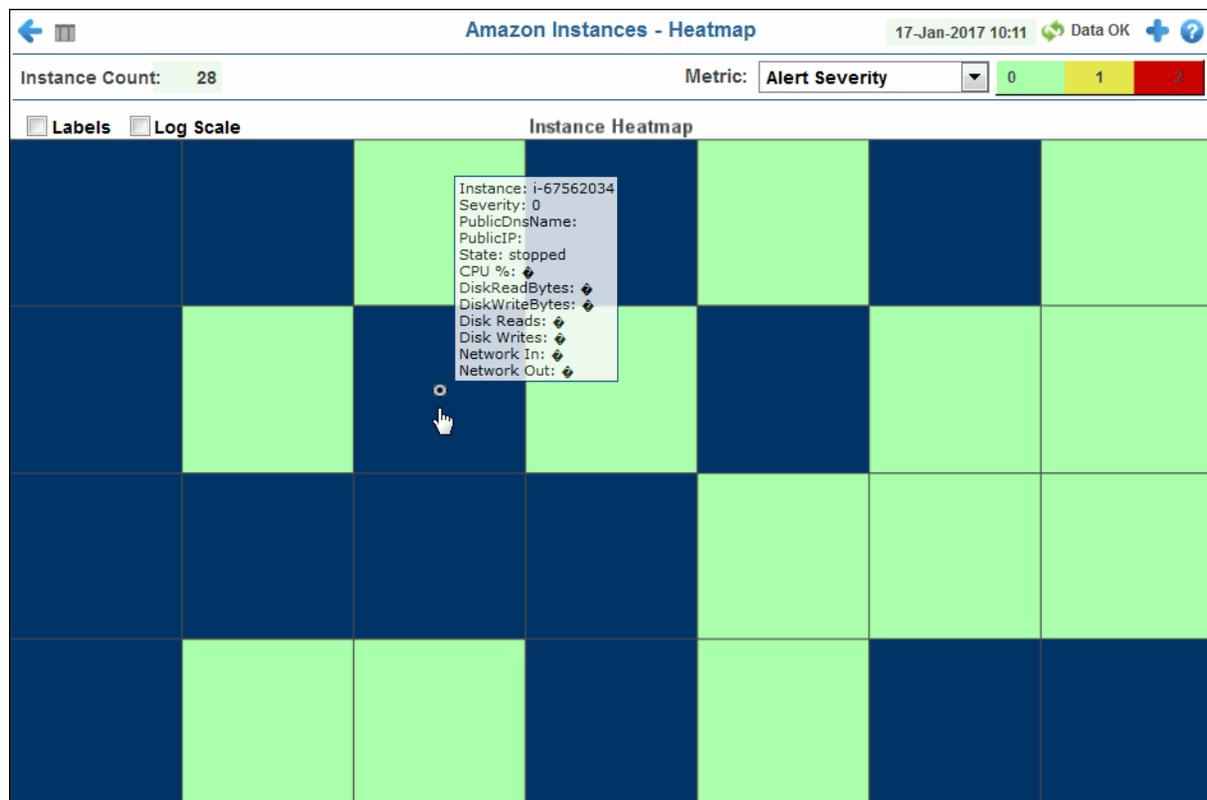
### Amazon EC2 Instance Heatmap

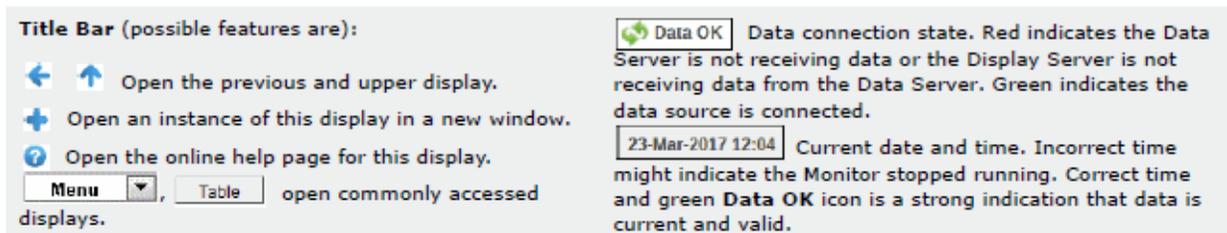
View the most critical alert states associated with your Amazon EC2 instances. Use this display to quickly identify instances with critical alerts. Compare heap usage, disk reads and writes and network throughput rates across all monitored instances.

Each rectangle in the heatmap represents an Amazon EC2 instance. The rectangle color indicates the most critical alert state associated with the instance for the selected **Metric**.

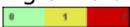
Choose a different metric to display from the **Metric** drop-down menu. Mouse over a rectangle to see additional metrics, including disk reads and writes, CPU utilization and network in/out rates. By default, this display shows **Alert Severity**.

Use the **Labels** check-box  to include or exclude labels in the heatmap. Click a rectangle to drill-down and view instance details in the ["Amazon EC2 Instance Summary"](#) display.





## Fields and Data:

<b>Instance Count:</b>	The total number of instances currently shown in the display.
<b>Labels:</b>	Select to show labels in the display.
<b>Log Scale</b>	Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.
<b>Metric</b>	Choose a metric to view in the display.
<b>Alert Severity</b>	<p>The maximum level of alerts in the rectangle. Values range from <b>0</b> - <b>2</b>, as indicated in the color gradient  bar, where <b>2</b> is the highest Alert Severity:</p> <ul style="list-style-type: none"> <li> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>CPU Usage</b>	The percent (%) CPU used. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum amount in the heatmap. The middle value in the gradient bar indicates the average amount.
<b>Network In</b>	The number of incoming bytes. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum amount in the heatmap. The middle value in the gradient bar indicates the average amount.
<b>Network Out</b>	The number of outgoing bytes. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum amount in the heatmap. The middle value in the gradient bar indicates the average amount.
<b>Disk Reads</b>	The number of completed disk reads. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum amount in the heatmap. The middle value in the gradient bar indicates the average amount.
<b>Disk Writes</b>	The number of completed disk writes. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum amount in the heatmap. The middle value in the gradient bar indicates the average amount.

**Disk Read Bytes**

The amount of disk reads, in bytes. The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum amount in the heatmap. The middle value in the gradient bar indicates the average amount.

**Disk Write Bytes**

The amount of disk writes, in bytes. The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum amount in the heatmap. The middle value in the gradient bar indicates the average amount.

**Amazon EC2 Instance Table**

View detailed utilization data for all your Amazon EC2 instances in a tabular format. Use this display to see all available data for this View.

Each row in the table is a different Amazon EC2 instance. Use the **Show:** drop-down menu to only show instances that are **running** or **stopped**.

Click a column header to sort column data in numerical or alphabetical order. Drill-down and investigate by clicking a row to view details for the selected application in the "Amazon EC2 Instance Summary" display.

Instance ID	Public DNS Name	State	% CPU	Disk Reads (bytes)	Disk Reads (Ops)	Disk Write (bytes)
i-094dec		stopped				
i-1a0a9e		stopped				
i-224dbe40		stopped				
i-2a31cdb3		stopped				
i-38405553		stopped				
i-3bd1fc50	ec2-75-101-146-223.compute-1.amazonaws.com	running	1.00	0	0	
i-41fa0fd8		stopped				
i-4a5ead28	ec2-54-144-4-125.compute-1.amazonaws.com	running	1.13	0	0	
i-570baf82		running	.40	0	0	
i-67562034		stopped				
i-6a3dd2f3		stopped				
i-6a838401	ec2-184-73-176-2.compute-1.amazonaws.com	running	98.67	0	0	
i-97280b20		stopped				
i-9a712903		running	.03	0	0	
i-9cac5205		stopped				
i-ac348fff		stopped				
i-aceeebc7	ec2-54-221-170-144.compute-1.amazonaws.com	running	1.27	0	0	
i-b6a6492f		stopped				
i-b7e3284d		running	4.57	0	0	
i-bbb3c1d1		stopped				
i-c5fad2ae	ec2-184-73-131-119.compute-1.amazonaws.com	running	1.13	0	0	
i-cfa13aab	ec2-54-235-4-76.compute-1.amazonaws.com	running	1.26	0	0	

**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** , **Table** open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Instance Count:** The number of instances in the table.

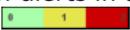
**Filter By:**

The display might include these filtering options:

**Show:** Choose to show **All** instances, **running** or **stopped** instances.

**All Instances Table:**

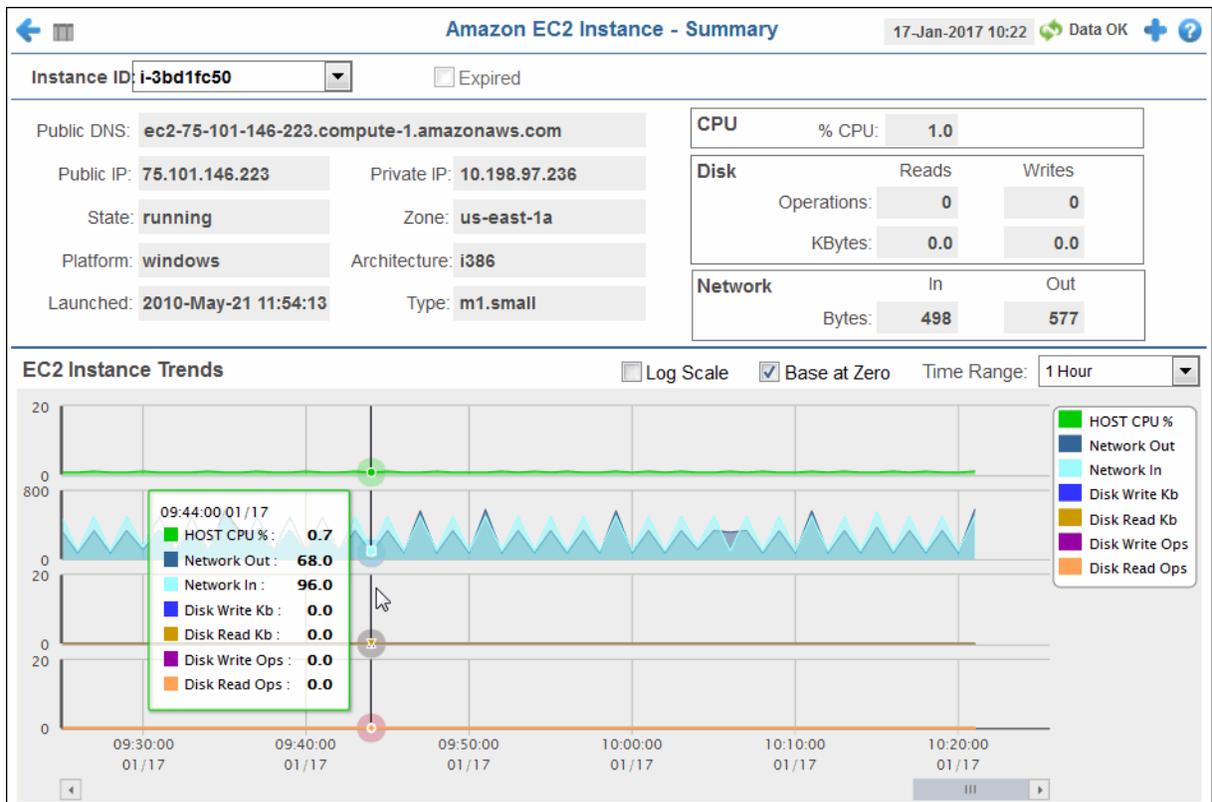
Each row in the table is a different instance.

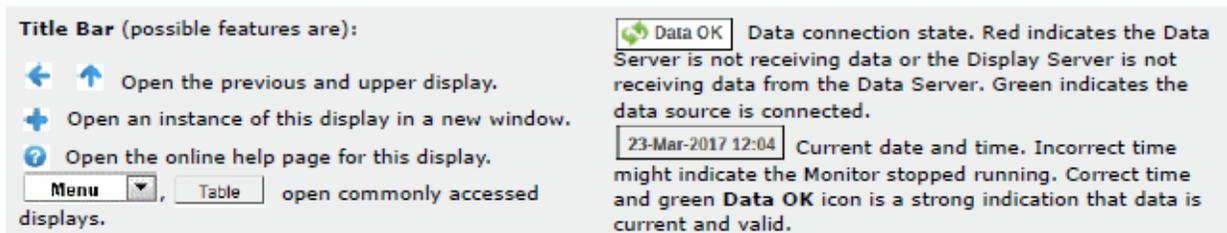
<b>Instance</b>	The name of the instance.
<b>Alert Severity</b>	The maximum level of alerts in the row. Values range from <b>0</b> - <b>2</b> , as indicated in the color gradient  bar, where <b>2</b> is the highest Alert Severity: <ul style="list-style-type: none"> <li><span style="color: red;">●</span> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li><span style="color: yellow;">●</span> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li><span style="color: green;">●</span> Green indicates that no metrics exceeded their alert thresholds.</li> </ul>
<b>Public DNS Name</b>	The public domain name of the instance.
<b>State</b>	The instance state ( <b>running</b> or <b>stopped</b> ).
<b>%CPU</b>	The percent CPU used.
<b>Disk Reads (bytes)</b>	The amount of disk reads, in bytes.
<b>Disk Reads (Ops)</b>	The number of disk reads (count).
<b>Disk Writes (bytes)</b>	The amount of disk writes, in bytes.
<b>Disk Writes (Ops)</b>	The number of disk writes (count).
<b>Network In</b>	The number of incoming bytes.
<b>Network Out</b>	The number of outgoing bytes.
<b>Instance Type</b>	The instance type (e.g. m1.small).
<b>Private IP</b>	The instance private IP address.
<b>Public IP</b>	The instance public IP address.
<b>Platform</b>	The instance operating system (e.g. windows).
<b>Architecture</b>	The instance architecture (e.g. i386).
<b>Image ID</b>	The unique identifier for the image. For details about Amazon EC2 data, refer to vendor documentation.
<b>Root Device Name</b>	The name of the root device. For details about Amazon EC2 data, refer to vendor documentation.
<b>Root Device Type</b>	The type of root device. For details about Amazon EC2 data, refer to vendor documentation.

- Availability Zone** The id for the availability zone (e.g. us-east-1a).  
For details about Amazon EC2 data, refer to vendor documentation.
- Group** For details, see vendor documentation.
- Tenancy** For details about Amazon EC2 data, refer to vendor documentation.
- Tags** For details, see vendor documentation.
- State Transition Reason** For details, see vendor documentation.
- LaunchTime** The date and time the instance was started.
- Timestamp** The date and time the data was last updated.
- Expired** When checked, data has not been received from this instance in the specified amount of time. The instance will be removed from the in the specified amount of time. The default setting is 60 seconds.

### Amazon EC2 Instance Summary

This display provides detailed utilization metrics for a single Amazon EC2 instance. Use this display to investigate performance details and trends for an instance.



**Filter By:**

The display might include these filtering options:

**Instance ID:** Choose an instance to show data for in the display.

**Fields and Data:**

Data describes the selected host except where noted.

	<b>Expired</b>	When checked, data has not been received from this instance in the specified amount of time. The instance will be removed from the in the specified amount of time. The default setting is 60 seconds.	
	<b>Public IP</b>	The instance public IP address.	
	<b>State</b>	The instance state ( <b>running</b> or <b>stopped</b> ).	
	<b>Platform</b>	The instance operating system (e.g. windows).	
	<b>Launched</b>	The date and time the instance was started.	
	<b>Private IP</b>	The instance private IP address.	
	<b>Zone</b>	The date and time the instance was started.	
	<b>Architecture</b>	The instance architecture (e.g. i386).	
	<b>Type</b>	The instance type (e.g. m1.small).	
<b>CPU</b>	<b>%CPU</b>	The percent CPU used.	
<b>Disk</b>	<b>Operations</b>	<b>Reads</b>	The number of disk reads (count).
		<b>Writes</b>	The number of disk writes (count).
	<b>KBytes</b>	<b>Reads</b>	The amount of disk reads, in kilobytes.
		<b>Writes</b>	The amount of disk writes, in kilobytes.
<b>Network</b>	<b>Bytes In</b>	The number of incoming bytes.	
	<b>Bytes Out</b>	The number of outgoing bytes.	

**EC2 Instance Trends**

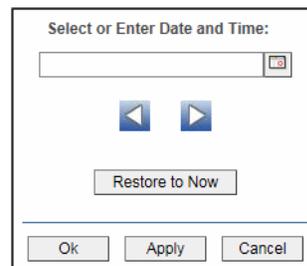
Traces metrics for the selected instance.

- **Host CPU%:** The amount of CPU used, in percent.
- **Network Out:** The number of outgoing bytes.
- **Network In:** The number of incoming bytes.
- **Disk Write Kb:** The amount of disk writes, in kilobytes.
- **Disk Read Kb:** The amount of disk reads, in kilobytes.
- **Disk Write Ops:** The number of disk writes (count).
- **Disk Read Ops:** The number of disk reads (count).

**Log Scale** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.

**Base at Zero** Select to use zero (**0**) as the Y axis minimum for all graph traces.

**Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

---

## Amazon Web Services - HTML

This section describes the HTML version of the Solution Package for Amazon Web Services.

The following Solution Package for Amazon Web Services HTML version features an overview display, "[Amazon EC2 Overview - HTML](#)" (shown below), and the following displays which can be found under **Components** tab > **Hosts** > **Amazon EC2 Hosts**. For additional details, see vendor documentation.

This section contains the following displays:

- ["Amazon EC2 Instance Heatmap - HTML"](#)
- ["Amazon EC2 Instance Table - HTML"](#)
- ["Amazon EC2 Instance Summary - HTML"](#)

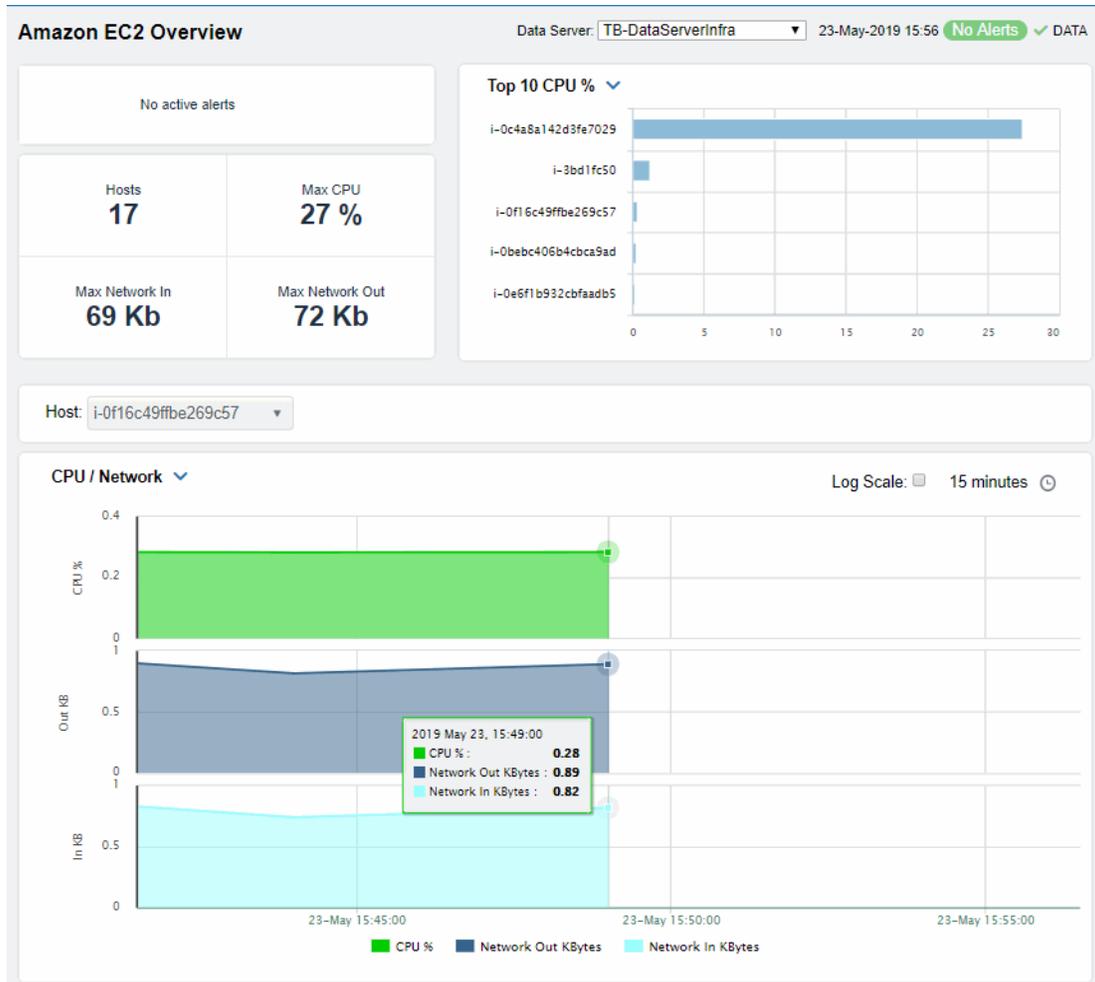
## **Amazon EC2 Overview - HTML**

The Amazon EC2 Overview is the top-level display for the Amazon Web Services Solution Package, which provides a good starting point for immediately getting the status of all your connections on your Data Server. You can select the RTView DataServer for which you want to see data and easily view the current data for that DataServer including:

- The total number of active alerts for the selected Data Server, including the total number of critical and warning alerts.
- The **Top 10 CPU%** hogs, **Network Out** (kilobytes sent) or **Network In** (kilobytes received) on your connected DataServer.

You can hover over each region in the upper half of the Overview to see more detail. You can also drill down to see even more detail by clicking on each respective region in the Overview.

The trend graph traces **CPU %**, **Network In** and **Network Out** for the selected **Host**. You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.



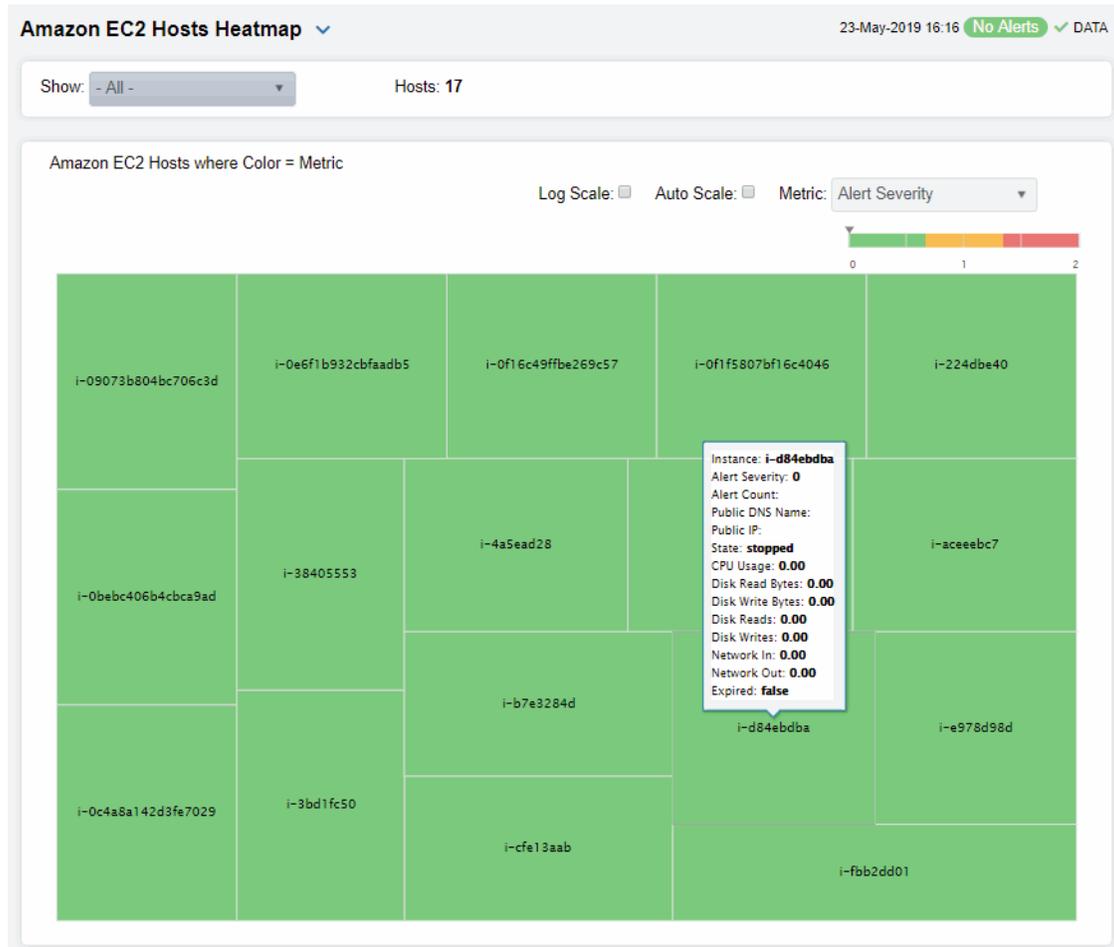
## Amazon EC2 Instance Heatmap - HTML

View the most critical alert states associated with your Amazon EC2 instances. Use this display to quickly identify instances with critical alerts. Compare heap usage, disk reads and writes and network throughput rates across all monitored instances.

Choose **All**, **Stopped** or **Running** hosts from the **Show** drop-down menu. Each rectangle in the heatmap represents an Amazon EC2 instance. The rectangle color indicates the most critical alert state associated with the instance for the selected **Metric**.

Choose a different metric to display from the **Metric** drop-down menu. Mouse over a rectangle to see additional metrics, including **Disk Reads** and **Writes**, **CPU Utilization** and **Network In/Out** rates. By default, this display shows **Alert Severity**.

Use the **Labels** check-box  to include or exclude labels in the heatmap. Click a rectangle to drill-down and view instance details in the ["Amazon EC2 Instance Summary - HTML"](#) display.



### Fields and Data:

**Instance Count:**

The total number of instances currently shown in the display.

**Labels:**

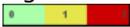
Select to show labels in the display.

**Log Scale**

Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.

**Metric**

Choose a metric to view in the display.

<b>Alert Severity</b>	<p>The maximum level of alerts in the rectangle. Values range from <b>0</b> - <b>2</b>, as indicated in the color gradient  bar, where <b>2</b> is the highest Alert Severity:</p> <ul style="list-style-type: none"> <li> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>CPU Usage</b>	<p>The percent (%) CPU used. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum amount in the heatmap. The middle value in the gradient bar indicates the average amount.</p>
<b>Network In</b>	<p>The number of incoming bytes. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum amount in the heatmap. The middle value in the gradient bar indicates the average amount.</p>
<b>Network Out</b>	<p>The number of outgoing bytes. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum amount in the heatmap. The middle value in the gradient bar indicates the average amount.</p>
<b>Disk Reads</b>	<p>The number of completed disk reads. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum amount in the heatmap. The middle value in the gradient bar indicates the average amount.</p>
<b>Disk Writes</b>	<p>The number of completed disk writes. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum amount in the heatmap. The middle value in the gradient bar indicates the average amount.</p>
<b>Disk Read Bytes</b>	<p>The amount of disk reads, in bytes. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum amount in the heatmap. The middle value in the gradient bar indicates the average amount.</p>
<b>Disk Write Bytes</b>	<p>The amount of disk writes, in bytes. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum amount in the heatmap. The middle value in the gradient bar indicates the average amount.</p>

## Amazon EC2 Instance Table - HTML

Investigate detailed configuration and utilization metrics for all Amazon EC2 instances. This display contains all metrics available for Amazon EC2 instances, including the **Public DNS Name**, **CPU%** utilization and **Alert Level**, where:

-  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
-  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
-  Green indicates that no metrics exceeded their alert thresholds.

Each row in the table contains data for a particular Amazon EC2 instance. You can search, filter, sort and choose columns to include by clicking a column header icon (to the right of each column label) and selecting Filter, Sort Ascending, Sort Descending or Columns. Or just click a column header to sort. Right-click on a table cell to **Export to Excel** or **Copy Cell Value**.

Use the **Show:** drop-down menu to only show instances that are **running** or **stopped**. Drill-down and investigate by clicking a row to view details for the selected instance in the Summary display.

Amazon EC2 Hosts Table 23-May-2019 16:19 No Alerts DATA

Show: - All - Hosts: 17

Tags	State Transition Reason	Launch Time	Time Stamp	Expired
e: Grahame - Red Hat 7.1	User initiated (2019-04-25 20:03:38 GMT)		23-May-2019 16:18:36	
e: Intuit WebLogic1	User initiated (2019-04-15 16:54:19 GMT)		23-May-2019 16:18:36	
e: KAFKA Server-1 and KAFKA Sender Docker	User initiated (2018-04-25 00:32:19 GMT)		23-May-2019 16:18:36	
e: Intuit WebLogic2	User initiated (2018-04-13 17:54:39 GMT)		23-May-2019 16:18:36	
e: Ed David SCAM Project	User initiated (2018-01-11 21:25:19 GMT)		23-May-2019 16:18:36	
e: Ed David SCAM 2 Project	User initiated (2017-10-05 00:20:28 GMT)		23-May-2019 16:18:36	
e: appmon-slcloud-clone1	User initiated (2017-05-26 18:55:13 GMT)		23-May-2019 16:18:36	
e: appmon.slcloud.com	Server.InternalError		23-May-2019 16:18:36	
e: appmon-slcloud-clone2	Server.InternalError		23-May-2019 16:18:36	
e: DOCS-SL-COM Apache SSL			23-May-2019 16:18:36	
e: DockerWorker1			23-May-2019 16:18:36	
e: RTView Automation Test - Muhammad			23-May-2019 16:18:36	
e: DockerWorker2			23-May-2019 16:18:36	
e: JBOSS			23-May-2019 16:18:36	
e: DayTrader			23-May-2019 16:18:36	
e: slidemos.com			23-May-2019 16:18:36	
e: JBOSS Rosario			23-May-2019 16:18:36	

**Instance Count:** The number of instances in the table.

**Filter By:**  
The display might include these filtering options:

**Show:** Choose to show **All** instances, **running** or **stopped** instances.

**All Instances Table:**  
Each row in the table is a different instance.

**Instance** The name of the instance.

**Alert Severity** The maximum level of alerts in the row. Values range from **0** - **2**, as indicated in the color gradient  bar, where **2** is the highest Alert Severity:

- Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
- Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
- Green indicates that no metrics exceeded their alert thresholds.

**Public DNS Name** The public domain name of the instance.

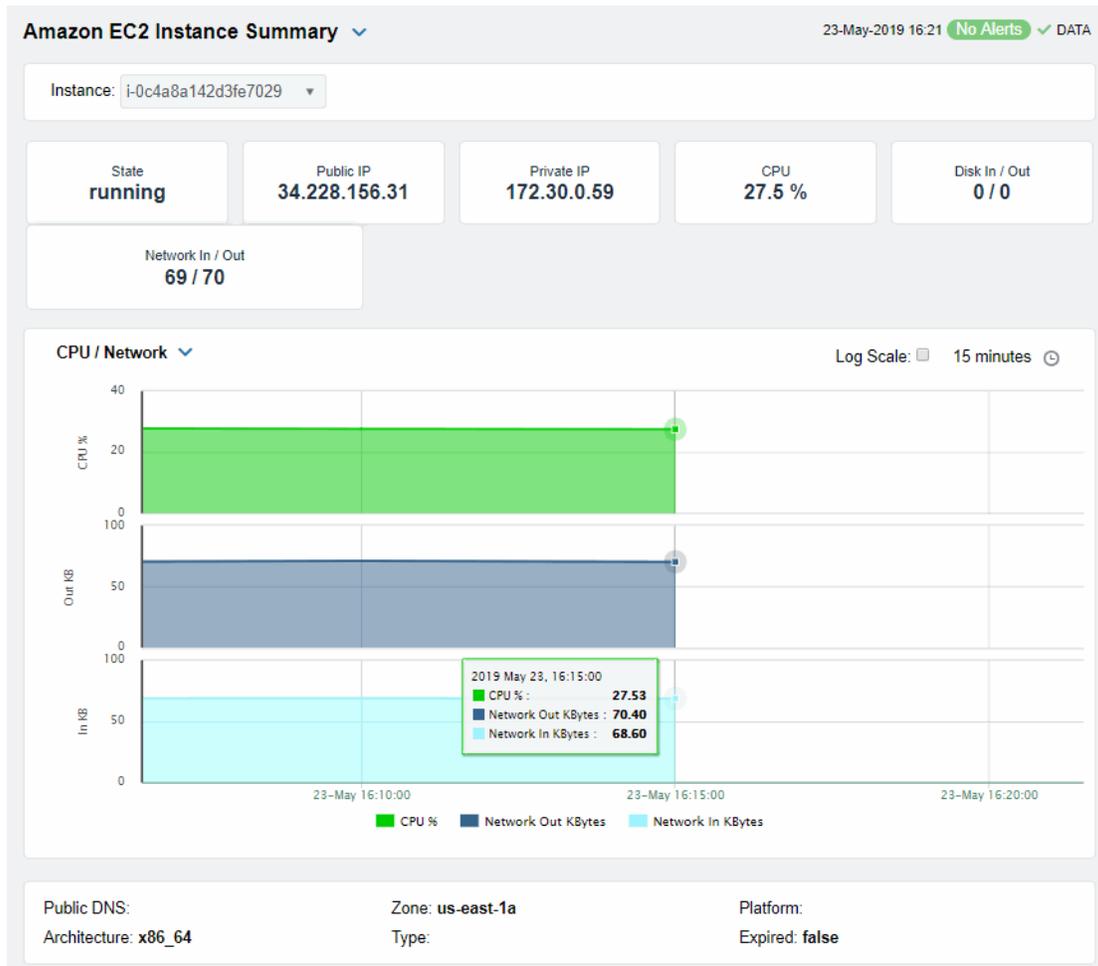
<b>State</b>	The instance state ( <b>running</b> or <b>stopped</b> ).
<b>%CPU</b>	The percent CPU used.
<b>Disk Reads (bytes)</b>	The amount of disk reads, in bytes.
<b>Disk Reads (Ops)</b>	The number of disk reads (count).
<b>Disk Writes (bytes)</b>	The amount of disk writes, in bytes.
<b>Disk Writes (Ops)</b>	The number of disk writes (count).
<b>Network In</b>	The number of incoming bytes.
<b>Network Out</b>	The number of outgoing bytes.
<b>Instance Type</b>	The instance type (e.g. m1.small).
<b>Private IP</b>	The instance private IP address.
<b>Public IP</b>	The instance public IP address.
<b>Platform</b>	The instance operating system (e.g. windows).
<b>Architecture</b>	The instance architecture (e.g. i386).
<b>Image ID</b>	The unique identifier for the image. For details about Amazon EC2 data, refer to vendor documentation.
<b>Root Device Name</b>	The name of the root device. For details about Amazon EC2 data, refer to vendor documentation.
<b>Root Device Type</b>	The type of root device. For details about Amazon EC2 data, refer to vendor documentation.
<b>Availability Zone</b>	The id for the availability zone (e.g. us-east-1a). For details about Amazon EC2 data, refer to vendor documentation.
<b>Group</b>	For details, see vendor documentation.
<b>Tenancy</b>	For details about Amazon EC2 data, refer to vendor documentation.
<b>Tags</b>	For details, see vendor documentation.
<b>State Transition Reason</b>	For details, see vendor documentation.
<b>LaunchTime</b>	The date and time the instance was started.
<b>Timestamp</b>	The date and time the data was last updated.
<b>Expired</b>	When checked, data has not been received from this instance in the specified amount of time. The instance will be removed from the in the specified amount of time. The default setting is 60 seconds.

## Amazon EC2 Instance Summary - HTML

Track utilization and performance metrics for a specific Amazon EC2 instance. Use this display to investigate performance details and trends for an instance.

You can choose to have the trend graph trace **CPU/Network** utilization and load values (**CPU %** utilization, **Network Out KBytes** and **Network In KBytes**), or **Disk** utilization and load values (**Disk Out KBytes**, **Disk In KBytes**, **Disk Out Ops** and **Disk In Ops**).

You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.



## Docker

The Solution Package for Docker is an easy to configure and use monitoring system that gives you extensive visibility into the health and performance of your Docker Engines, Docker Containers, and the applications that rely on them.

The following Docker Views can be found under **Components** tab > **Processes** > **Docker Engines**:

- **"Engine View"**: The displays in this View allow you to view the current and historical metrics for all engines in a heatmap or tabular format for one or all hosts, or view the current and historical metrics for a single engine.
- **"Container View"**: The displays in this View allow you to view the current and historical metrics for all containers in a heatmap or tabular format for one or all hosts, or view the current and historical metrics for a single container.

## Engine View

These displays provide detailed data for all engines or for a particular engine. Displays in this View are:

- **"Engines Heatmap"**: A heatmap view of all engines and their associated metrics.
- **"Engines Table"**: A tabular view of your engines and their associated metrics.
- **"Engine Summary"**: Provides additional details and a way to view trending data for a single engine.

## Engines Heatmap

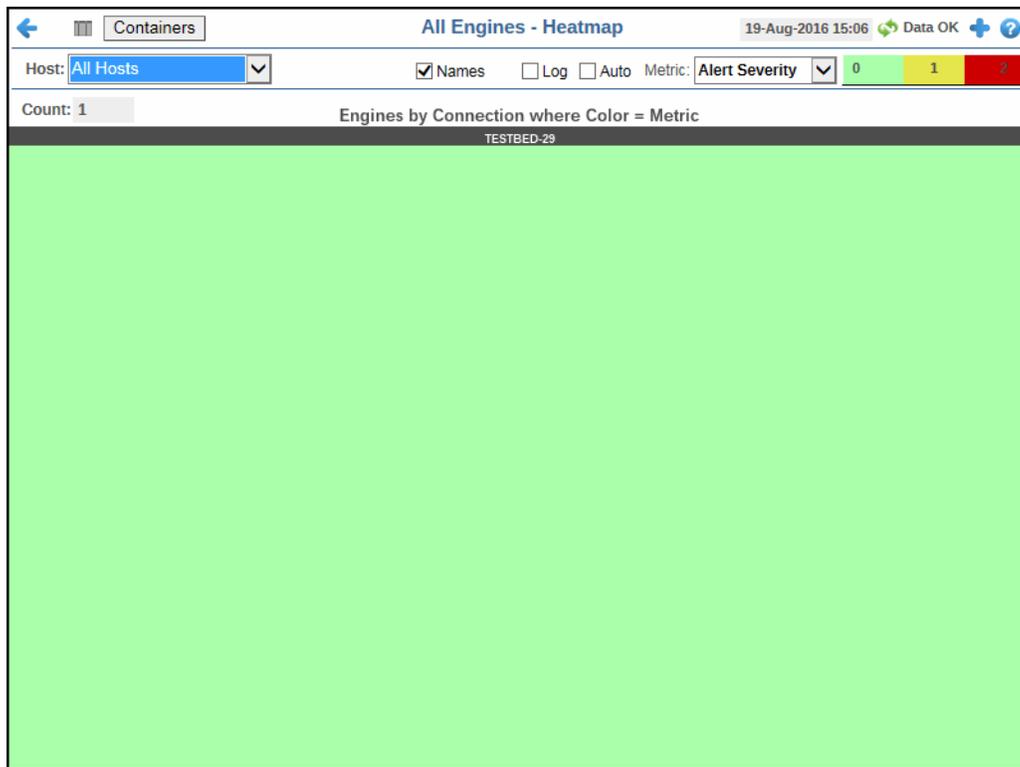
This heatmap display provides an easy-to-view interface that allows you to quickly identify the current status of each of your engines for each available metric. You can view the engines in the heatmap based on the following metrics: the current alert severity, the current alert count, the percentage of CPU used, the amount of memory used, the total incoming bytes, and the total outgoing bytes. By default, this display shows the heatmap based on the **Alert Severity** metric.

You can use the **Names** check-box  to include or exclude labels in the heatmap, and you can mouse over a rectangle to see additional metrics for an engine. Clicking one of the rectangles in the heatmap opens the **"Engine Summary"** display, which allows you to see additional details for the selected engine.

---

**Note:** When the data for the engine being monitored expires, the color of the rectangle representing that engine in the heatmap automatically changes to a color that is not included in the color gradient bar so that you can easily identify when the data is stale. Expired data could occur for a number of reasons including, but not limited to, the connection to the engine may have been lost, or the engine could have experienced a problem and may no longer be up-and-running.

---



#### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- The number of items currently in the display.

- Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

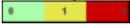
#### Fields and Data:

- Host** Select the host for which you want to show data in the display.
- Count** Lists the total number of engines found using the search parameters.
- Names** Select this check box to display the names of the engines at the top of each rectangle in the heatmap.
- Log** Select this check box to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Auto** Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value.  
**Note:** Some metrics auto-scale automatically, even when **Auto** is not selected.

**Metric**

Choose a metric to view in the display.

**Alert Severity**

The current alert severity. Values range from **0** - **2**, as indicated in the color gradient  bar, where **2** is the highest Alert Severity:

-  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
-  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
-  Green indicates that no metrics have exceeded their alert thresholds.

**Alert Count**

The total number of critical and warning unacknowledged alerts in the engine. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.

**CPU Usage**

The percentage of CPU used by the engine. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **DocEngineCpuUsageHigh**. The middle value in the gradient bar indicates the middle value of the range.

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

**Memory**

The current memory usage by the engine, in kilobytes, which includes all memory regardless of when it was accessed. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of connections in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

The **Auto** option does not impact this metric.

**Net Bytes In**

The total number of incoming bytes. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **DocEngineNetBytesInHigh**. The middle value in the gradient bar indicates the middle value of the range.

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

**Net Bytes Out**

The total number of outgoing bytes. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **DocEngineNetBytesOutHigh**. The middle value in the gradient bar indicates the middle value of the range.

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

## Engines Table

This table provides a view of all of your engines and their associated metric data including host, alert severity, alert count, and the current value of each gathered metric. You can click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by clicking a row to view details for the selected engine in the “Engine Summary” display

Host	Alert Level	Alert Count	CPU Usage	Memory Available (KB)	Memory Usage (KB)	Memory WS (KB)	Memory RSS (KB)	Memory Limited	Net Bytes In avg
TESTBED-29		0	11.39	3,782,232	3,373,604	1,602,908	58,564	<input checked="" type="checkbox"/>	81,200

### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- The number of items currently in the display.

- Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

---

**Note:** The **Containers** button takes you to “Containers Table”.

---

### Fields and Data:

#### Host

Select the name of the host (or **All Hosts**) containing the engines for which you want to view data.

**Count** The total number of engines being monitored based on your search criteria.

**All Engines Table:**

**Host** The name of the host.

**Alert Level** The current alert severity.  
● Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  
● Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  
● Green indicates that no metrics have exceeded their alert thresholds.

**Alert Count** The total number of alerts for the host.

**CPU Usage** The percentage of CPU used by the engine.

**Memory Available (KB)** The amount of memory, in kilobytes, that is available to the engine.

**Memory Usage (KB)** The current memory usage by the engine, in kilobytes, which includes all memory regardless of when it was accessed.

**Memory WS (KB)** The amount of memory (in kilobytes) in the working set, which includes recently accessed memory, dirty memory, and kernel memory.

**Memory RSS (KB)** The amount of anonymous and swap cache memory (including transparent/hugepages), in kilobytes.

**Memory Limited** When checked, the amount of memory available to the engine is limited.

**Net Bytes In avg** The average number of incoming bytes per second.

**Net Bytes Out avg** The average number of outgoing bytes per second.

**Net Packets In avg** The average number of incoming packets per second.

**Net Packets Out avg** The average number of outgoing packets per second.

**Docker Version** The Docker software version of the Docker Engine.

**Container OS Version** The version of the container’s operating system on which the docker engine is running.

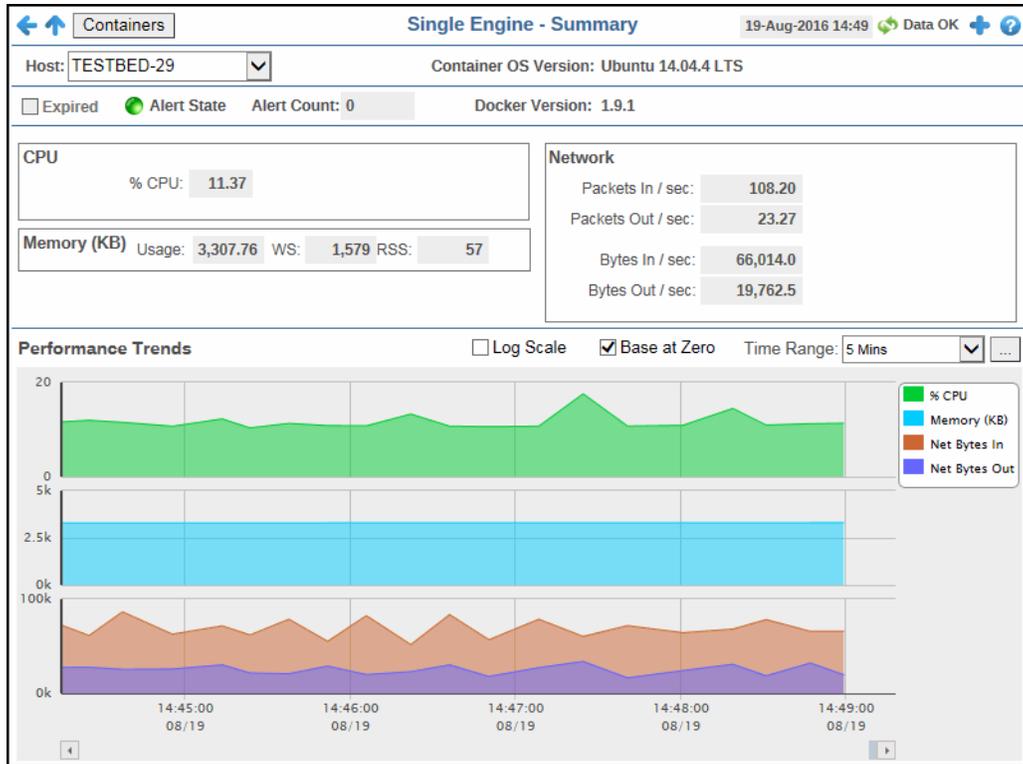
**Container Kernal Version** The version of the container’s Kernal in which the docker engine is running.

**Expired** When checked, performance data has not been received within the time specified (in seconds) in the **Expire Time** field in the **Duration** region in the RTView Configuration Application > (**Project Name**) > **Solution Package Configuration** > **Docker** > **DATA STORAGE** tab. The **Delete Time** field (also in the **Duration** region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.

**Timestamp** The date and time the row data was last updated.

## Engine Summary

This display allows you to view current as well as trending data for the percentage of CPU used by the engine, memory usage details, and network data details.



**Title Bar (possible features are):**

- Open the previous and upper display.
- Open the online help page for this display.
- open commonly accessed displays.
- 6,047** The number of items currently in the display.
- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

**Note:** The **Containers** button takes you to "Containers Table".

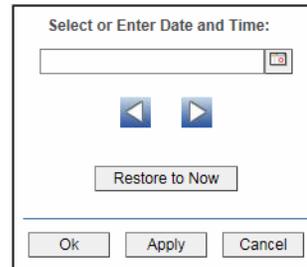
### Filter By:

- Host** Select the host for which you want to show data in the display.
- Container OS Version** The version of the container's operating system on which the docker engine is running.

**Fields and Data:**

<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > ( <b>Project Name</b> ) > <b>Solution Package Configuration</b> > <b>Docker</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Alert State</b>	The current alert severity. <ul style="list-style-type: none"> <li><span style="color: red;">●</span> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li><span style="color: yellow;">●</span> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li><span style="color: green;">●</span> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Alert Count</b>	The total number of current alerts.
<b>Docker Version</b>	The Docker software version of the Docker Engine.
<b>CPU</b>	
<b>% CPU</b>	The percentage of CPU used by the engine.
<b>Memory (KB)</b>	
<b>Usage</b>	The current memory usage by the engine, in kilobytes, which includes all memory regardless of when it was accessed.
<b>WS</b>	The amount of memory (in kilobytes) in the working set, which includes recently accessed memory, dirty memory, and kernel memory.
<b>RSS</b>	The Resident Set Size, which is the amount of anonymous and swap cache memory (including transparent/hugepages), in kilobytes.
<b>Network</b>	
<b>Packets In/sec</b>	The average number of incoming packets per second..
<b>Packets Out/sec</b>	The average number of outgoing packets per second.
<b>Bytes In/sec</b>	The average number of incoming bytes per second.
<b>Bytes Out/sec</b>	The average number of outgoing bytes per second.
<b>Performance Trends Graph</b>	Traces the following: <ul style="list-style-type: none"> <li><b>% CPU</b> -- traces the percentage of CPU being used on the engine.</li> <li><b>Memory (KB)</b> -- traces the amount of memory, in kilobytes, used by the engine.</li> <li><b>Net Bytes In</b> -- traces the average number of incoming bytes per second.</li> <li><b>Net Bytes Out</b> -- traces the average number of outgoing bytes per second.</li> </ul> <p><b>Log Scale</b>      Select to enable a logarithmic scale. Use <b>Log Scale</b> to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. <b>Log Scale</b> makes data on both scales visible by applying logarithmic values rather than actual values to the data.</p>

- Base at Zero** Select to use zero (**0**) as the Y axis minimum for all graph traces.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Container View

These displays allow you to view the current and historical metrics for all containers in a heatmap or tabular format for one or all hosts, or view the current and historical metrics for a single container. Displays in this View are:

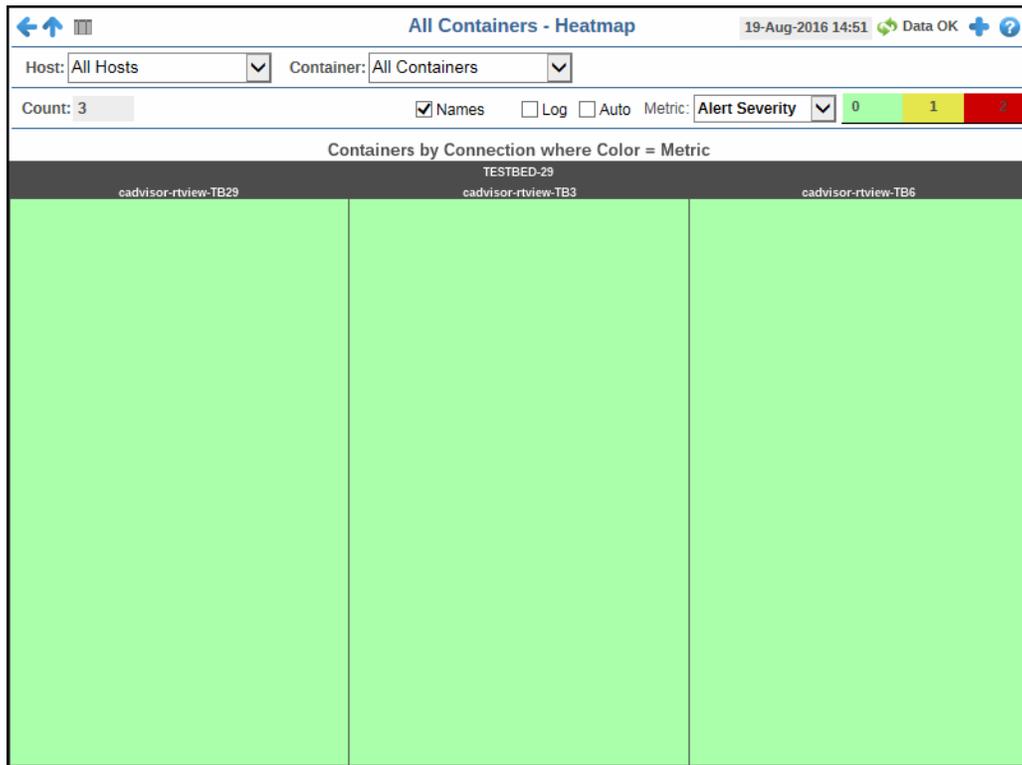
- **"Containers Heatmap"**: A color-coded heatmap view of data for all containers for a particular host.
- **"Containers Table"**: A tabular view of data for all containers for a particular host.
- **"Container Summary"**: This display allows you to view current and trending data for a single container for a particular host.

## Containers Heatmap

This heatmap display provides an easy-to-view interface that allows you to quickly identify the current status of each of your containers for each available metric. You can view the containers in the heatmap based on the following metrics: the current alert severity, the current alert count, the percentage of CPU used, and the percentage of memory used. By default, this display shows the heatmap based on the **Alert Severity** metric.

You can use the **Names** check-box  to include or exclude labels in the heatmap, and you can mouse over a rectangle to see additional metrics for a container. Clicking one of the rectangles in the heatmap opens the “**Container Summary**” display, which allows you to see additional details for the selected container.

**Note:** When the data for the container being monitored expires, the color of the rectangle representing that container in the heatmap automatically changes to a color that is not included in the color gradient bar so that you can easily identify when the data is stale. Expired data could occur for a number of reasons including, but not limited to, the connection to the container may have been lost, or the container could have experienced a problem and may no longer be up-and-running.

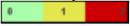
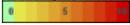


**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- The number of items currently in the display.
- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

**Fields and Data:**

**Host** Select the host (or **All Hosts**) for which you want to show data in the heatmap.

<b>Container</b>	Select the container (or <b>All Containers</b> ) for which you want to show data in the heatmap..
<b>Count</b>	Lists the total number of containers (rows) found using the search parameters.
<b>Names</b>	Select this check box to display the names of the containers at the top of each rectangle in the heatmap.
<b>Log</b>	Select this check box to enable a logarithmic scale. Use <b>Log Scale</b> to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. <b>Log Scale</b> makes data on both scales visible by applying logarithmic values rather than actual values to the data.
<b>Auto</b>	Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value. <b>Note:</b> Some metrics auto-scale automatically, even when <b>Auto</b> is not selected.
<b>Metric</b>	Choose a metric to view in the display.
<b>Alert Severity</b>	<p>The current alert severity. Values range from <b>0</b> - <b>2</b>, as indicated in the color gradient  bar, where <b>2</b> is the highest Alert Severity:</p> <ul style="list-style-type: none"> <li> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Alert Count</b>	The total number of critical and warning unacknowledged alerts in the instance. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.
<b>CPU Usage</b>	<p>The percentage of CPU used by the container. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of <b>DocContainerCpuUsageHigh</b>. The middle value in the gradient bar indicates the middle value of the range.</p> <p>When <b>Auto</b> is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>
<b>Memory</b>	<p>The current memory usage by the container, in kilobytes, which includes all memory regardless of when it was accessed. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of connections in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p> <p>The <b>Auto</b> option does not impact this metric.</p>

**Net Bytes In** The number of incoming bytes per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **DocContainerNetBytesInHigh**. The middle value in the gradient bar indicates the middle value of the range.

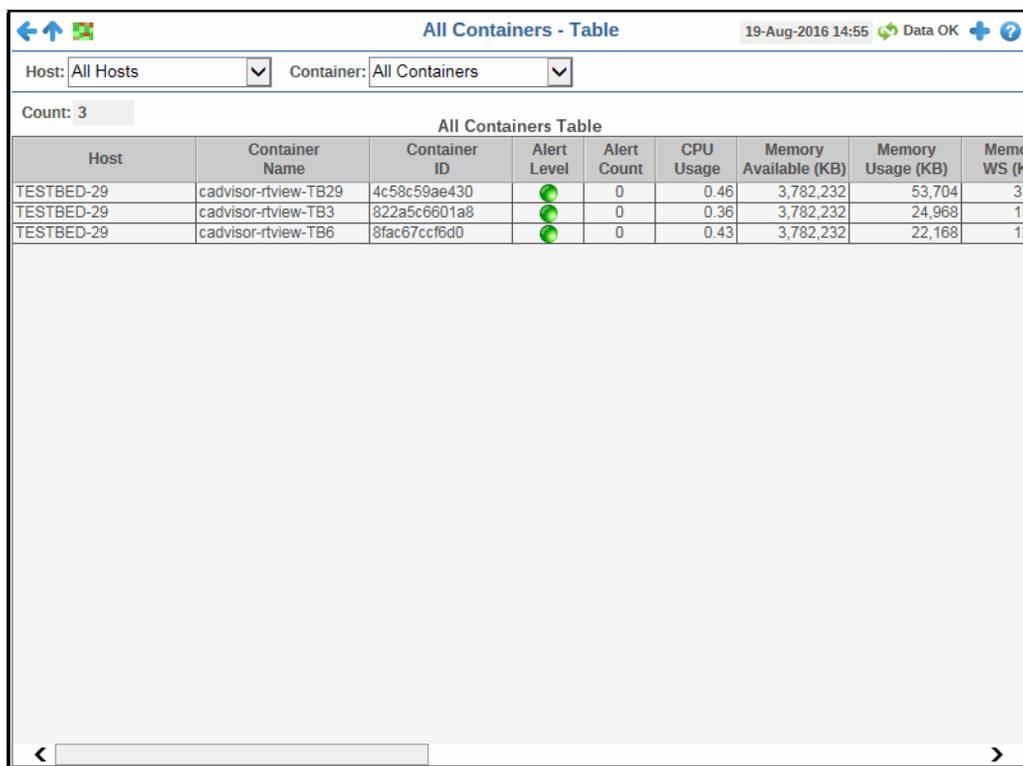
When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

**Net Bytes Out** The number of outgoing bytes per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **DocContainerNetBytesOutHigh**. The middle value in the gradient bar indicates the middle value of the range.

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

## Containers Table

This display allows you to view details in a table format for one container on a particular host, for all containers on a particular host, for a particular container on all hosts, or for all containers on all hosts. You can drill-down and view the details for a particular container in the "Container Summary" display by clicking on a row in the resulting table.

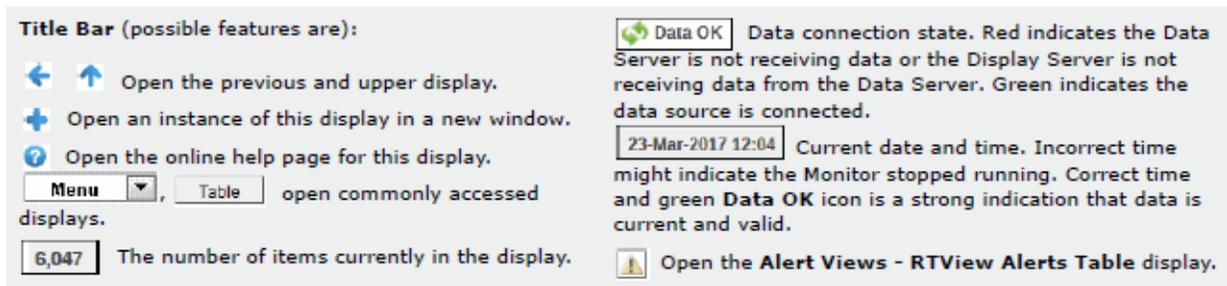


Host: All Hosts Container: All Containers

Count: 3

All Containers Table

Host	Container Name	Container ID	Alert Level	Alert Count	CPU Usage	Memory Available (KB)	Memory Usage (KB)	Memory WS (KB)
TESTBED-29	cadvisor-rtview-TB29	4c58c59ae430		0	0.46	3,782,232	53,704	3
TESTBED-29	cadvisor-rtview-TB3	822a5c6601a8		0	0.36	3,782,232	24,968	1
TESTBED-29	cadvisor-rtview-TB6	8fac67ccf6d0		0	0.43	3,782,232	22,168	1

**Filter By:**

The display includes these filtering options:

- Host** Select the host for which you want to show data in the display.
- Container** Select the container (or **All Containers**) for which you want to view data..
- Count** Lists the total number of containers (rows) found using the search parameters.

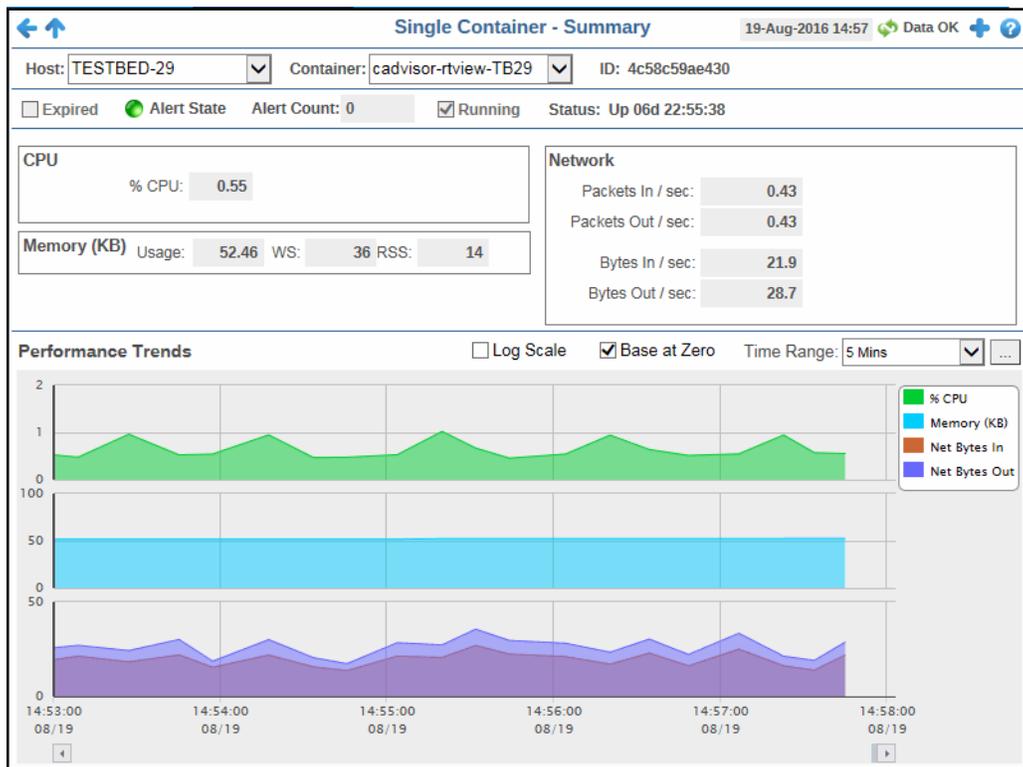
**All Containers Table**

- Host** The name of the host.
- Container Name** The name of the container.
- Container ID** The absolute container name.
- Alert Level** The current alert status.
  - Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
  - Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
  - Green indicates that no metrics have exceeded their alert thresholds.
- Alert Count** Total number of alerts for the process.
- CPU Usage** The percentage of CPU used by the container.
- Memory Available (KB)** The amount of memory, in kilobytes, that is available to the container.
- Memory Usage (KB)** Current memory usage by the container, in kilobytes, which includes all memory regardless of when it was accessed.
- Memory WS (KB)** The amount of memory (in kilobytes) in the working set, which includes recently accessed memory, dirty memory, and kernel memory.
- Memory RSS (KB)** The Resident Set Size, which is the amount of anonymous and swap cache memory (including transparent/hugepages), in kilobytes.
- Memory Limited** When checked, the amount of memory available to the container is limited. If not checked, then the amount of memory available to the container is unlimited, which means the amount of memory available to the container is the same as the memory available to the engine.
- Net Bytes In avg** The average number of incoming bytes per second.
- Net Bytes Out avg** The average number of outgoing bytes per second.

<b>Net Packets In avg</b>	The average number of incoming packets per second.
<b>Net Packets Out avg</b>	The average number of outgoing packets per second.
<b>Uptime</b>	The amount of time (in seconds) that the container has been up and running.
<b>Running</b>	When checked, this check box indicates that the container is running.
<b>Status</b>	<p>The current status of the container. Values are:</p> <p><b>Up</b> - indicates that the container is up and running, and lists the amount of time the container has been up and running (<b>Uptime</b>).</p> <p><b>Created</b> - indicates that the container has been created but is currently not in use.</p> <p><b>Exited</b> - indicates that the container has been stopped, and lists the error code as well as the amount of time since the container was stopped.</p>
<b>Starts</b>	<p>The number of times the container (re)started within the time specified (in seconds) in the <b>\$docEventCacheTimeRange</b> field in the <b>conf\rtvapm_dockermon.properties</b> file. The default is 3600 seconds (1 hour). For example, by default, this row column lists the number of times the container has (re)started in the past hour. This number provides a good indication of the stability of the container; the higher the number, the more unstable the container.</p>
<b>Expired</b>	<p>When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application &gt; (<b>Project Name</b>) &gt; <b>Solution Package Configuration</b> &gt; <b>Docker</b> &gt; <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.</p>
<b>Timestamp</b>	The date and time the row data was last updated.

## Container Summary

This display provides a view of the current and historical metrics for a single container. You can view the current information pertaining to CPU usage percentage, Memory details, Disk read and write details, and network data details in the upper portion of the display. The trend graph in the bottom half of the display traces the current and historical CPU usage, the average memory used, and the number of incoming and outgoing network bytes.



**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** , **Table** open commonly accessed displays.
- 6,047** The number of items currently in the display.
- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

**Filter By:**

The display might include these filtering options:

- Host** Select the host for which you want to show data in the display.
- Container** Select the container for which you want to show data in the display.
- ID** The absolute container name.

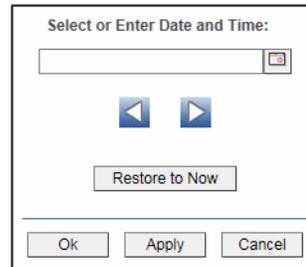
**Fields and Data:**

**Expired** When checked, performance data has not been received within the time specified (in seconds) in the **Expire Time** field in the **Duration** region in the RTView Configuration Application > **(Project Name)** > **Solution Package Configuration** > **Docker** > **DATA STORAGE** tab. The **Delete Time** field (also in the **Duration** region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.

<b>Alert State</b>	<p>The current alert severity.</p> <ul style="list-style-type: none"> <li><span style="color: red;">●</span> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li><span style="color: yellow;">●</span> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li><span style="color: green;">●</span> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Alert Count</b>	The total number of current alerts.
<b>Running</b>	When checked, this check box indicates that the container is running.
<b>Status</b>	<p>The current status of the container. Values are:</p> <p><b>Up</b> - indicates that the container is up and running, and lists the amount of time the container has been up and running (<b>Uptime</b>).</p> <p><b>Created</b> - indicates that the container has been created but is currently not in use.</p> <p><b>Exited</b> - indicates that the container has been stopped, and lists the error code as well as the amount of time since the container was stopped.</p>
<b>CPU</b>	
<b>% CPU</b>	The percentage of CPU used by the container.
<b>Memory (KB)</b>	
<b>Usage</b>	The current memory usage by the container, in kilobytes, which includes all memory regardless of when it was accessed.
<b>WS</b>	The amount of memory (in kilobytes) in the working set, which includes recently accessed memory, dirty memory, and kernel memory.
<b>RSS</b>	The Resident Set Size, which is the amount of anonymous and swap cache memory (including transparent/hugepages), in kilobytes.
<b>Network</b>	
<b>Packets In/sec</b>	The average number of incoming packets per second.
<b>Packets Out/sec</b>	The average number of outgoing packets per second.
<b>Bytes In/sec</b>	The average number of incoming bytes per second.
<b>Bytes Out/sec</b>	The average number of outgoing bytes per second.
<b>Performance Trends Graph</b>	<p>Traces the following:</p> <ul style="list-style-type: none"> <li><b>% CPU</b> -- traces percentage of CPU used by the container.</li> <li><b>Memory (KB)</b> -- traces the current memory usage by the container, in kilobytes, which includes all memory regardless of when it was accessed.</li> <li><b>Net Bytes In</b> -- traces the average number of incoming bytes per second.</li> <li><b>Net Bytes Out</b> -- traces the average number of outgoing bytes per second.</li> </ul> <p><b>Log Scale</b> Select to enable a logarithmic scale. Use <b>Log Scale</b> to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. <b>Log Scale</b> makes data on both scales visible by applying logarithmic values rather than actual values to the data.</p>

**Base at Zero** Select to use zero (**0**) as the Y axis minimum for all graph traces.

**Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

---

## Docker - HTML

The following Docker Views and their associated displays are available in the Monitor. This section describes the Monitor displays and includes:

- ["Docker Overview - HTML"](#): Describes the **Docker Overview** display.
- ["Docker Engines View - HTML"](#): The displays in this View allow you to view the current and historical metrics for all engines in a heatmap or tabular format for one or all hosts, or view the current and historical metrics for a single engine.
- ["Docker Containers View - HTML"](#): The displays in this View allow you to view the current and historical metrics for all containers in a heatmap or tabular format for one or all hosts, or view the current and historical metrics for a single container.

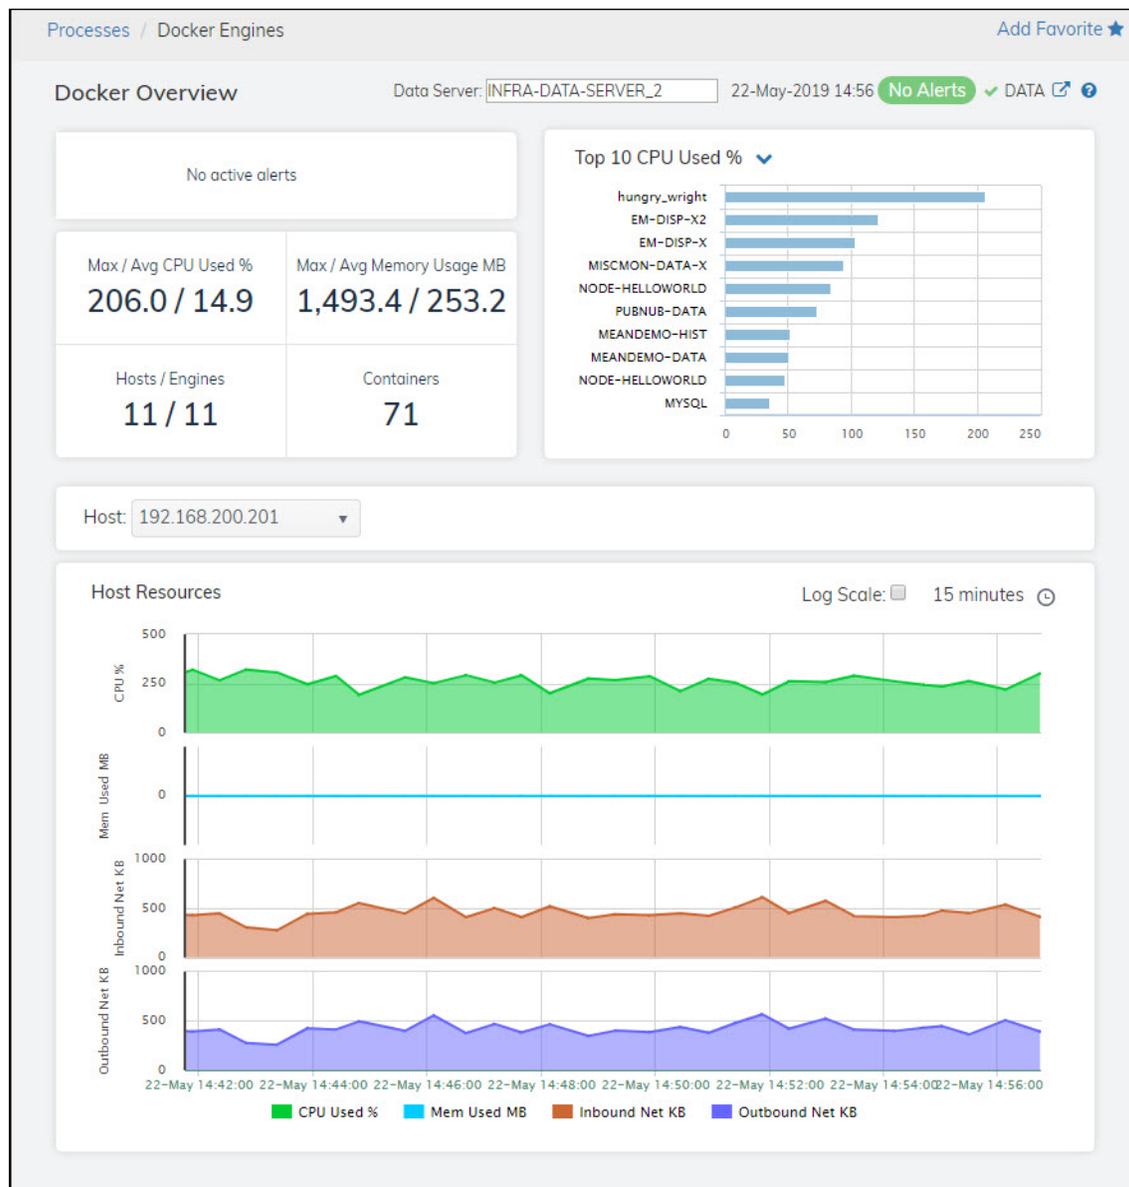
## Docker Overview - HTML

The **Docker Overview** is the top-level display for the Docker Monitor, which provides a good starting point for immediately getting the status of all your engines and containers on your Data Server. You can select the RTView DataServer for which you want to see data and easily view the current data for that DataServer including:

- The total number of active alerts for the selected DataServer, including the total number of critical and warning alerts.
- The maximum and average CPU percentage used across all containers.
- The maximum and average memory usage (in megabytes) across all containers.
- The total number of running hosts and the total number of engines.
- The total number of containers on your connected DataServer.
- A visual list of the top 10 containers based on CPU used percentage, memory used, inbound net kilobytes, and outbound net kilobytes on your connected DataServer.

You can hover over each region in the upper half of the Overview to see more detail. You can also drill down to see even more detail by clicking on each respective region in the Overview. For example, clicking on the alerts in the **CRITICAL** and **WARNING** alerts region opens the **Alerts Table by Components** display.

The bottom half of the display provides a host resources trend graph for a selected host. You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.



## Docker Engines View - HTML

These displays provide detailed data for all engines or for a particular engine. Clicking **Docker Engines** from the left/navigation menu opens the "[Docker Engines Table - HTML](#)" display, which shows a tabular view of all engines and their associated metrics. The options available under **Docker Engines** are:

- **All Engines Heatmap**: Opens the "[Docker Engines Heatmap - HTML](#)" display, which provides a heatmap view of all engines and their associated metrics.
- **Single Engine Summary**: Opens the "[Docker Single Engine Summary - HTML](#)" display, which provides additional details and a way to view trending data for a single engine.

## Docker Engines Table - HTML

This table provides a view of all of your engines and their associated metric data including host, alert severity, alert count, and the current value of each gathered metric. Each row in the table contains data for a particular engine. Click a column header to sort column data in ascending or descending order. Double-click on a table row to drill-down to the ["Docker Single Engine Summary - HTML"](#) display and view metrics for that particular engine. Toggle between the commonly accessed displays by clicking the drop down list on the display title.

Docker Engines Table 22-May-2019 15:24 No Alerts [DATA](#) [?](#)

Engines: 11

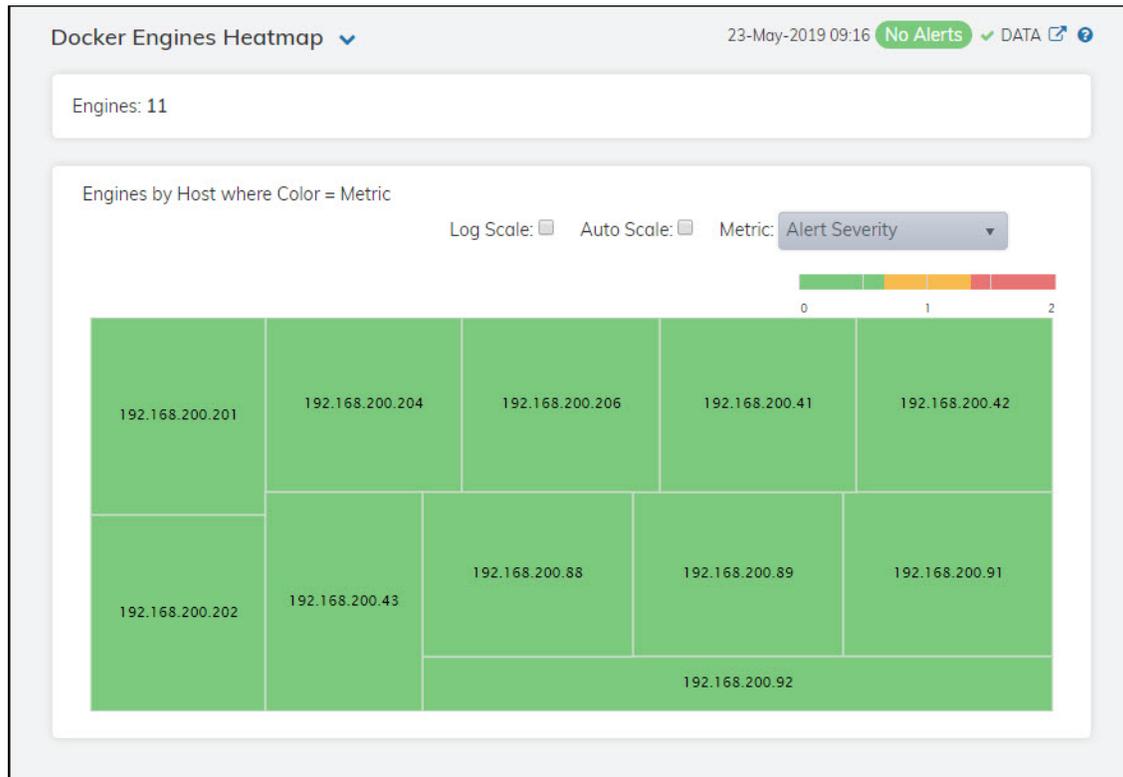
Host	Alert Level	Alert Count	CPU Usage %	Memory Avg MB	Memory Use MB	Memory WS MB	Memory RSS MB
192.168.200.201	✓		285.73	0.0	3,153.6	2,735.5	205.1
192.168.200.202	✓		59.05	0.0	1,317.2	1,120.5	178.1
192.168.200.204	✓		5.78	0.0	1,037.1	484.2	55.1
192.168.200.206	✓		7.53	0.0	909.4	598.4	277.1
192.168.200.41	✓		2.83	0.0	3,136.5	2,817.0	1,047.1
192.168.200.42	✓		98.84	0.0	3,100.9	1,388.1	673.1
192.168.200.43	✓		10.81	0.0	3,140.7	1,831.5	1,670.1
192.168.200.88	✓		2.98	0.0	1,507.1	1,210.1	425.1
192.168.200.89	✓		58.3	0.0	1,106.4	1,020.1	189.1
192.168.200.91	✓		194.47	0.0	5,121.4	2,363.5	288.1
192.168.200.92	✓		166.91	0.0	5,270.0	2,629.1	290.1

## Docker Engines Heatmap - HTML

Clicking **All Engines Heatmap** in the left/navigation menu opens the **Docker Engines Heatmap**, which provides an easy-to-view interface that allows you to quickly identify the current status of each of your engines for each available metric. You can view the engines in the heatmap based on the following metrics: the current alert severity, the current alert count, the percentage of CPU used, the amount of memory used, the total incoming bytes, and the total outgoing bytes. By default, this display shows the heatmap based on the **Alert Severity** metric.

Each rectangle in the heatmap represents an engine. The rectangle color indicates the most critical alert state associated with the engine. Choose a different metric to display from the **Metric** drop-down menu. Mouse over a rectangle to see additional metrics. By default, this display shows **Alert Severity**.

Drill-down and investigate an engine by clicking a rectangle in the heatmap to view details in the ["Docker Single Engine Summary - HTML"](#) display.



### Available Metrics

Select the metric driving the heatmap display. The default is **Alert Severity**. Each **Metric** has a color gradient bar that maps values to colors. The heatmap organizes the servers by host, where each rectangle represents an engine. Mouse-over any rectangle to display the current values of the metrics for the engine. Click on a rectangle to drill-down to the associated ["Docker Single Engine Summary - HTML"](#) display for a detailed view of metrics for that particular engine.

#### Alert Severity

The current alert severity. Values range from **0** - **2**, as indicated in the color gradient bar, where **2** is the highest Alert Severity:

- Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
- Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
- Green indicates that no metrics have exceeded their alert thresholds.

#### Alert Count

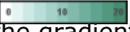
The total number of critical and warning unacknowledged alerts in the engine. The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.

**CPU Usage**

The percentage of CPU used by the engine. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **DocEngineCpuUsageHigh**. The middle value in the gradient bar indicates the middle value of the range.

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

**Memory Usage MB**

The current memory usage by the engine, in megabytes, which includes all memory regardless of when it was accessed. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of connections in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

The **Auto** option does not impact this metric.

**Net Inbound Avg KB**

The average number of incoming kilobytes. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **DocEngineNetBytesInHigh**. The middle value in the gradient bar indicates the middle value of the range.

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

**Net Outbound Avg KB**

The average number of outgoing kilobytes. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **DocEngineNetBytesOutHigh**. The middle value in the gradient bar indicates the middle value of the range.

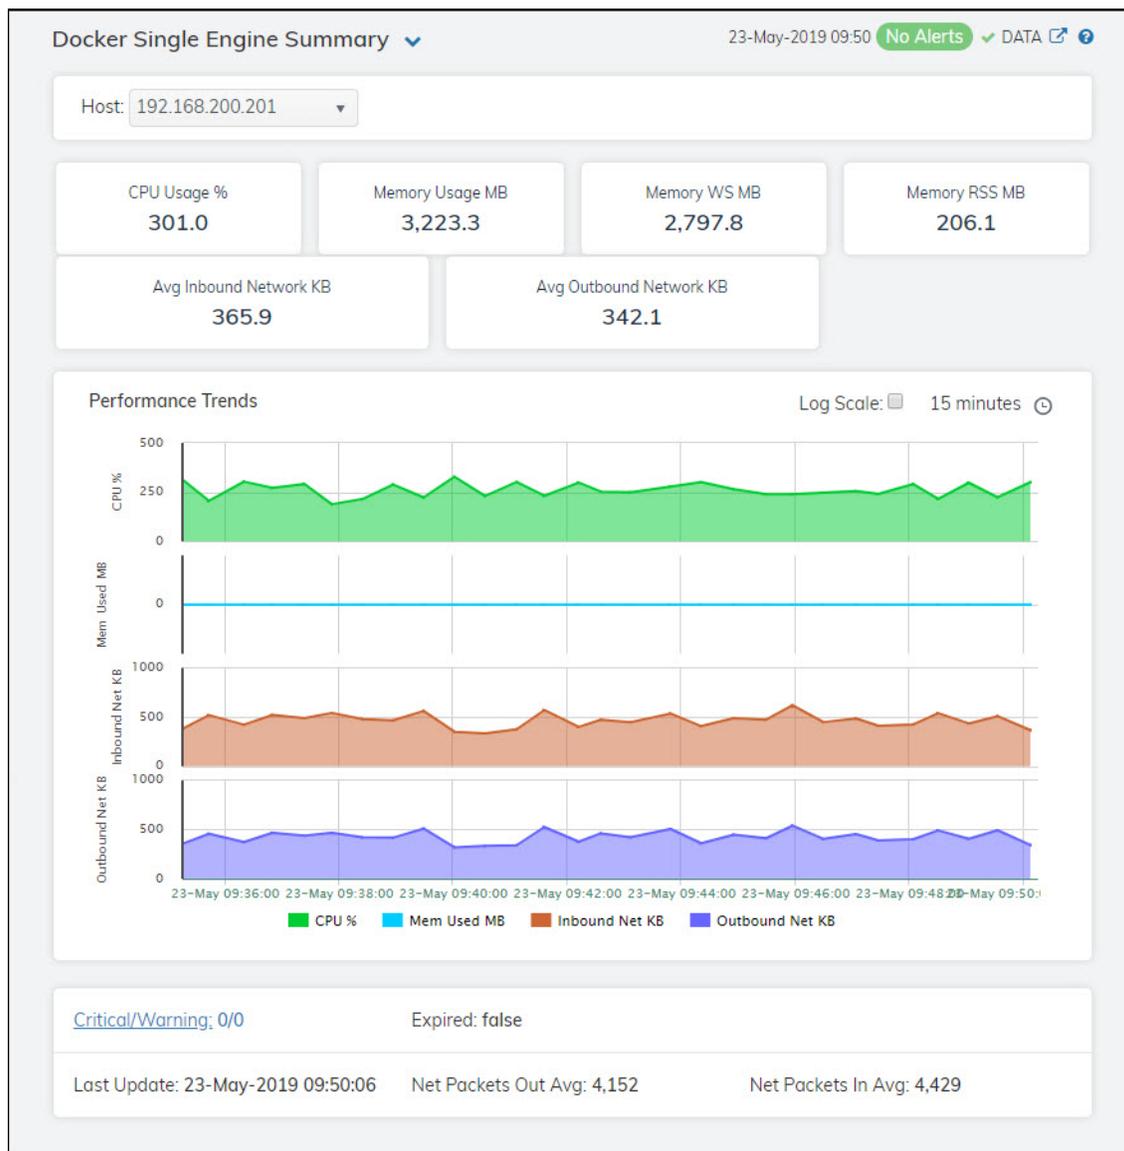
When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

## Docker Single Engine Summary - HTML

Clicking **Single Engine Summary** in the left/navigation menu opens the **Docker Single Engine Summary** display, which allows you to view current as well as trending data for the percentage of CPU used by the engine, memory usage details, and network data details. Clicking on the information boxes at the top of the display takes you to the "[Docker Engines Table - HTML](#)" display, where you can view additional engines data.

The **Performance Trends** trend graph allows you to view trend data for the CPU percentage, memory used, inbound net kilobytes, and outbound net kilobytes over a selected time range.

Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display.



## Docker Containers View - HTML

These displays allow you to view the current and historical metrics for all containers in a heatmap or tabular format for one or all hosts, or view the current and historical metrics for a single container. Clicking **Docker Containers** from the left/navigation menu opens the ["Docker Containers Table - HTML"](#) display, which shows a tabular view of all containers and their associated metrics for a particular host. The options available under **Docker Containers** are:

- **All Containers Heatmap:** Opens the ["Docker Containers Heatmap - HTML"](#) display, which provides a heatmap view of all containers and their associated metrics for a particular host.
- **Single Container Summary:** Opens the ["Docker Single Container Summary- HTML"](#) display, which provides additional details and a way to view trending data for a single container for a particular host.

## Docker Containers Table - HTML

This display allows you to view details in a table format for all containers on a particular host or for all containers on all hosts. You can drill-down and view the details for a particular container in the ["Docker Single Container Summary- HTML"](#) display by double-clicking on a row in the resulting table.

The screenshot shows a web interface titled "Docker Containers Table". At the top right, it displays the date and time "23-May-2019 11:10", a "No Alerts" status with a green checkmark, and icons for "DATA", a link, and a help icon. Below the title, there is a dropdown menu for "Host" currently set to "192.168.200.201". Underneath, it says "Containers: 11". The main part of the interface is a table titled "Containers Table".

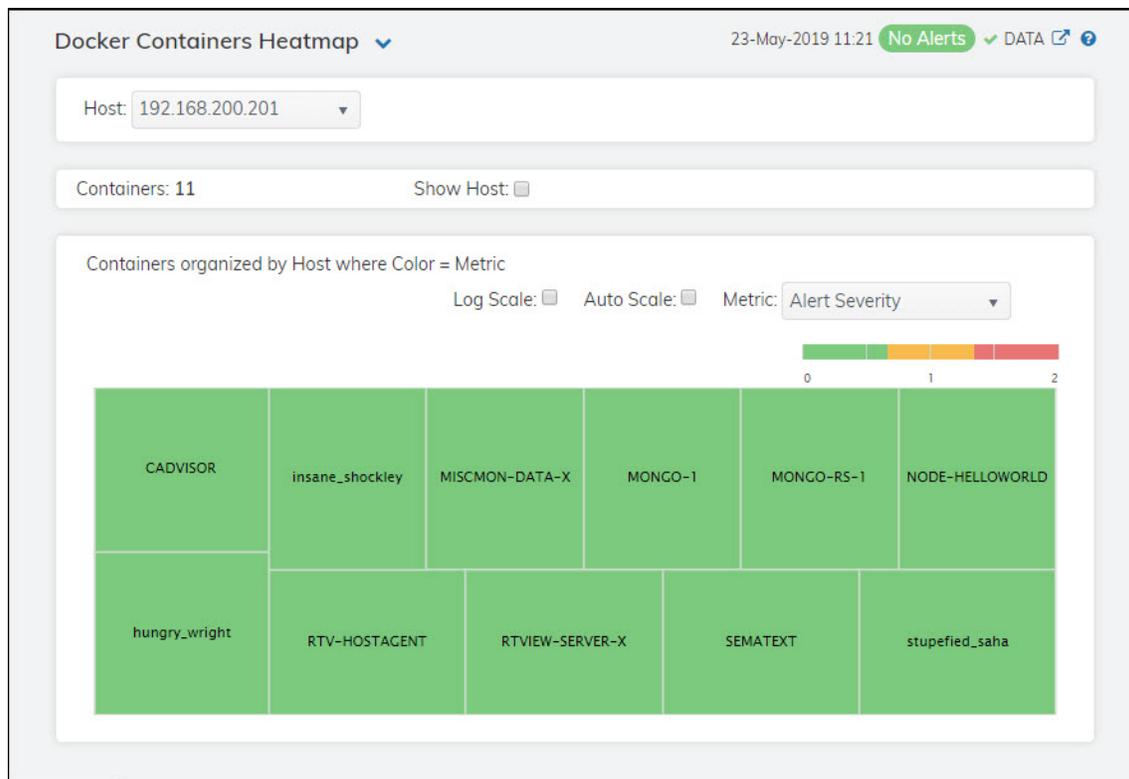
Host	Container	ID	Alert Level	Alert Count	CPU Usage %	Memory Avg MB
192.168.200.201	CADVISOR	c11f6657c487a	✓		3.05	
192.168.200.201	hungry_wright	08e8e8258671f	✓		193.49	
192.168.200.201	insane_shockley	d7dfd0f9cc01a0	✓		0.78	
192.168.200.201	MISCMON-DATA-X	53c89d6f7f350	✓		93.66	
192.168.200.201	MONGO-1	439618f9f2c16	✓		0.47	
192.168.200.201	MONGO-RS-1	fd28c47fb254e	✓		1.17	
192.168.200.201	NODE-HELLOWORLD	7ac31b3c23f5e	✓		99.7	
192.168.200.201	RTV-HOSTAGENT	1880c92f87f14	✓		0.06	
192.168.200.201	RTVIEW-SERVER-X	1914cb45ca69b	✓		0.31	
192.168.200.201	SEMATEXT	b1eb5a090454	✓		1.03	
192.168.200.201	stupefied_saha	9fbad85f9f428	✓		0.79	

## Docker Containers Heatmap - HTML

Clicking **All Containers Heatmap** in the left/navigation menu opens the **Docker Containers Heatmap**, which provides an easy-to-view interface that allows you to quickly identify the current status of each of your containers for each available metric. You can view the containers in the heatmap based on the following metrics: the current alert severity, the current alert count, the percentage of CPU used, and the percentage of memory used. By default, this display shows the heatmap based on the **Alert Severity** metric.

Each rectangle in the heatmap represents a container. The rectangle color indicates the most critical alert state associated with the container. Choose a different metric to display from the **Metric** drop-down menu. You can use the **Show Hosts** check-box  to include or exclude labels in the heatmap, and you can mouse over a rectangle to see additional metrics for a container.

Drill-down and investigate an engine by clicking a rectangle in the heatmap to view details in the ["Docker Single Container Summary- HTML"](#) display.



### Available Metrics

Select the metric driving the heatmap display. The default is **Alert Severity**. Each **Metric** has a color gradient bar that maps values to colors. The heatmap organizes the servers by host, where each rectangle represents a container. Mouse-over any rectangle to display the current values of the metrics for the container. Click on a rectangle to drill-down to the associated ["Docker Single Container Summary- HTML"](#) display for a detailed view of metrics for that particular container.

**Alert Severity** The current alert severity. Values range from **0** - **2**, as indicated in the color gradient  bar, where **2** is the highest Alert Severity:

-  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
-  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
-  Green indicates that no metrics have exceeded their alert thresholds.

**Alert Count** The total number of critical and warning unacknowledged alerts in the instance. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.

**CPU Usage %** The percentage of CPU used by the container. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **DocContainerCpuUsageHigh**. The middle value in the gradient bar indicates the middle value of the range.

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

**Memory Usage MB** The current memory usage by the container, in megabytes, which includes all memory regardless of when it was accessed. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of connections in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

The **Auto** option does not impact this metric.

**Net Avg Bytes In** The net average number of incoming bytes. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **DocContainerNetBytesInHigh**. The middle value in the gradient bar indicates the middle value of the range.

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

**Net Avg Bytes Out** The net average number of outgoing bytes. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **DocContainerNetBytesOutHigh**. The middle value in the gradient bar indicates the middle value of the range.

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

## Docker Single Container Summary- HTML

Clicking **Single Container Summary** in the left/navigation menu opens the **Docker Single Container Summary** display, which provides a view of the current and historical metrics for a single container. You can view the current information pertaining to CPU usage percentage, Memory details, and network data details in the upper portion of the display. Clicking on the information boxes at the top of the display takes you to the ["Docker Containers Table - HTML"](#) display, where you can view additional containers data.

The **Performance Trends** trend graph allows you to view trend data for the CPU percentage, memory used, inbound net kilobytes, and outbound net kilobytes over a selected time range.

Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display.



## JBoss

The following JBoss Views can be found under **Components** tab > **Application/Web Servers**> **JBoss**:

- **"JBoss Servers View"**: The displays in this View present server performance metrics such as CPU and memory utilization.
- **"JBoss Applications"**: The displays in this View present views of the VPN-level metrics.

With the Solution Package for Red Hat® JBoss® you are able to identify potential problems before they become critical and impact overall application performance. Typical installations of RTView Enterprise and its solution packages take only a few hours, while developing custom views for a variety of IT and development roles can be achieved in just days.



## JBoss Servers View

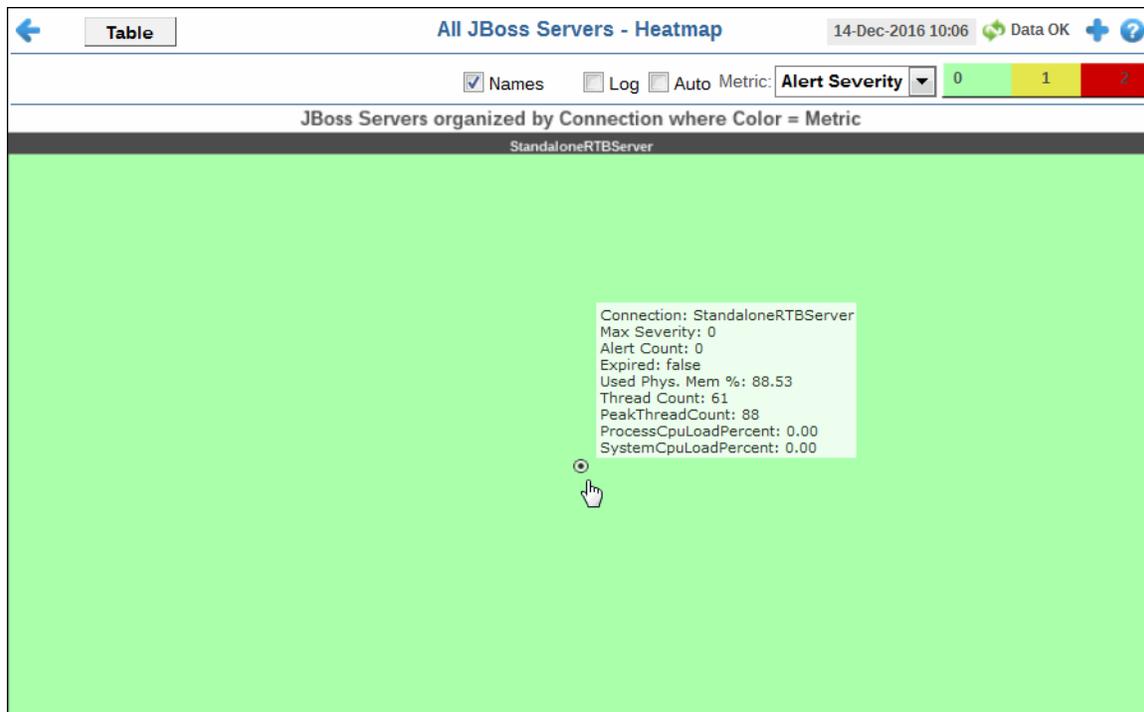
Displays in this View are:

- **"All Servers Heatmap"**
- **"All Servers Table"**
- **"Server Summary"**

## All Servers Heatmap

This heatmap shows the current status of connections on all JBoss servers. Use this display to quickly assess the current status of connections using various metrics, including **Alert Count** and **CPU Used %** and **Virtual Memory Used %**. By default, this display shows the heatmap based on the **Alert Severity** metric.

Each rectangle is a different JBoss server. Use the **Names** check-box  to include or exclude labels in the heatmap, and mouse over a rectangle to see additional metrics for a server. Click a rectangles to drill down to the "Server Summary" display, which shows additional details about the server.



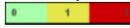
### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- 6,047** The number of items currently in the display.

- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

### Fields and Data:

- Names** Select to include labels in the heatmap.
- Log** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.

<b>Auto</b>	Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value. <b>Note:</b> Some metrics auto-scale automatically, even when <b>Auto</b> is not selected.
<b>Metric</b>	Choose a metric to view in the display.
<b>Alert Severity</b>	The current alert severity. Values range from <b>0</b> - <b>2</b> , as indicated in the color gradient  bar, where <b>2</b> is the highest Alert Severity: <ul style="list-style-type: none"> <li> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Alert Count</b>	The total number of critical and warning unacknowledged alerts. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.
<b>CPU Used%</b>	The percent CPU used. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of <b>JbsServerCpuUsedHigh</b> . The middle value in the gradient bar indicates the middle value of the range.  When <b>Auto</b> is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.
<b>V Memory Used%</b>	The percent virtual memory used. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of <b>JbsServerMemUsedHigh</b> . The middle value in the gradient bar indicates the middle value of the range.  When <b>Auto</b> is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.
<b>Free Memory</b>	The total amount of available memory. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum amount of available memory. The middle value in the gradient bar indicates the average amount.  The <b>Auto</b> flag does not have any impact on this metric.

## All Servers Table

View JBoss server details per connection such as the total number of sessions, bytes sent/received, and processing time. Each row in the table is a different server. The row color for inactive servers is dark red.

Drill-down and investigate by clicking a row in the table to view details for the selected connection in the “[Server Summary](#)” display.

Connection	Name	Expired	Alert Level	Process Cpu %	Threads	Peak Threads	Daemon Threads
StandaloneRTBServer	1996@WIN-8-CLONE	<input type="checkbox"/>	<span style="color: green;">●</span>	0.00	58	88	39

#### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

#### Fields and Data

This display includes:

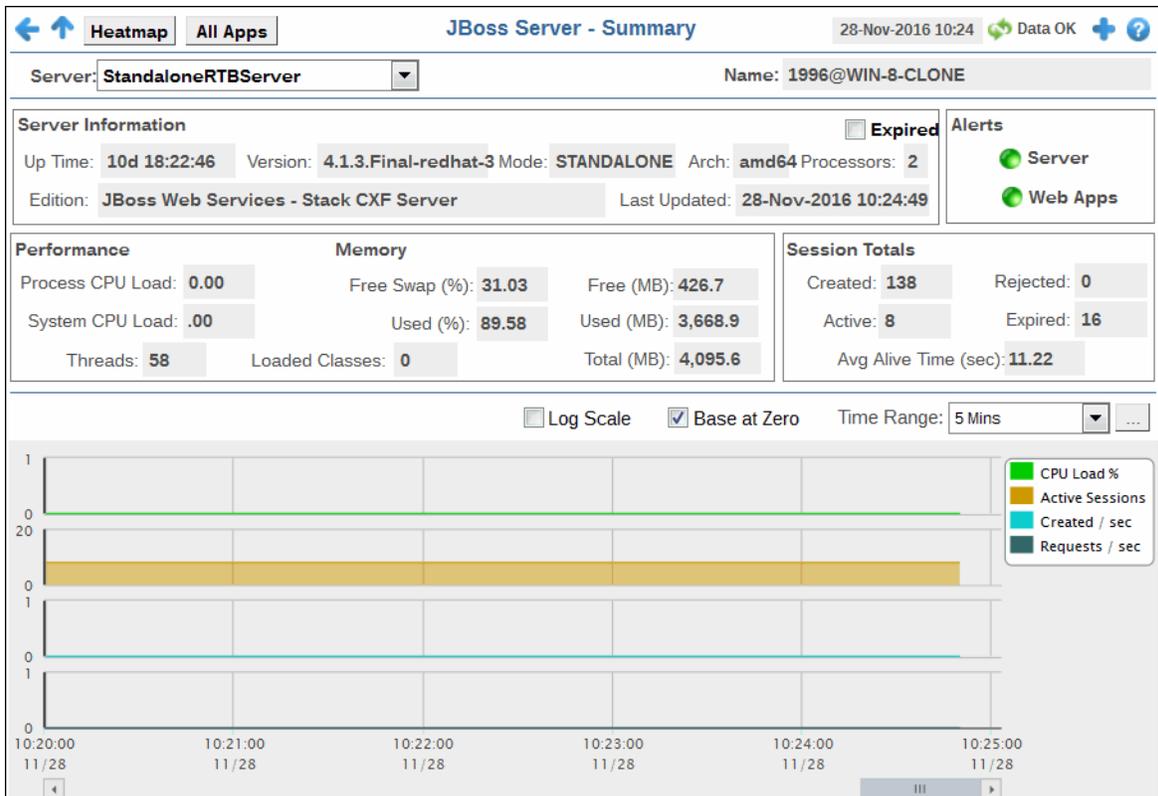
- Connection** The name of the connected server.
- Name** The name of the connection.
- Expired** When checked, data has not been received from this host in the specified amount of time. The host will be removed from the in the specified amount of time. The default setting is **60** seconds.
- Alert Level** The current alert severity.
  - Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
  - Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
  - Green indicates that no metrics have exceeded their alert thresholds.
- Process Cpu%** The amount of CPU used by processes, in percent.
- Threads** The total number of currently active threads.
- Peak Threads** The maximum number of active threads.

**Daemon Threads** The total number of currently active daemon threads.

**Started Threads** The total number of threads started since the server was last started.

## Server Summary

Track the performance of one server, see detailed current information as well as historical trends. You can drill down to this display from the ["All Servers Heatmap"](#) and ["All Servers Table"](#) displays.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu**, **Table** open commonly accessed displays.
- 6,047** The number of items currently in the display.

- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

### Fields and Data

This display includes:

**Server** Select a server.

**Name** The name of the connection.

### Server Information

**Up Time** The amount of time that the server has been up and running.

**Version** The version of the operating system.

**Mode** The current server mode:

- STANDALONE
- DOMAIN MODE

**Arch** The type of server architecture.

**Processors** The number of processors on the server.

**Expired** When checked, data has not been received from this server in the specified amount of time. The server will be removed from the per the specified amount of time. The default setting is **35** seconds.

**Edition** Refers to the vendor edition.

**Last Updated** The date and time of the last data update.

**Alerts** The current alert severity.

 Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.

 Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.

 Green indicates that no metrics have exceeded their alert thresholds.

### Performance

**Process CPU Load** The amount of process CPU utilization, in megabytes.

**System CPU Load** The amount of process CPU utilization, in megabytes.

**Threads** The number of active threads.

**Load Classes** The number of active load classes.

### Memory

**Free Swap %** The amount of free swap memory, in percent.

**Free MB** The amount of free swap memory, in megabytes.

**Used %** The amount of used memory, in percent.

**Used MB** The amount of used memory, in megabytes.

**Total MB** The memory sum total (**Free MB + Used MB**), in megabytes.

### Session Totals

**Created** The total number of sessions created since the server was restarted.

**Rejected** The total number of sessions rejected since the server was restarted.

**Active** The total number of currently active sessions.

**Expired** The total number of currently expired sessions.

**Average Alive Time sec** The average amount of time per session, in seconds.

**Trend Graph**

Traces metrics for the selected server.

- Log Scale** Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Base at Zero** Use zero as the Y axis minimum for all graph traces.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

- CPU Load %** Traces the amount of process CPU utilization, in percent.
- Active Sessions** Traces the number of active sessions.
- Created/sec** Traces the number of active sessions per second.
- Requests/sec** Traces the number of requests per second.

**JBoss Applications**

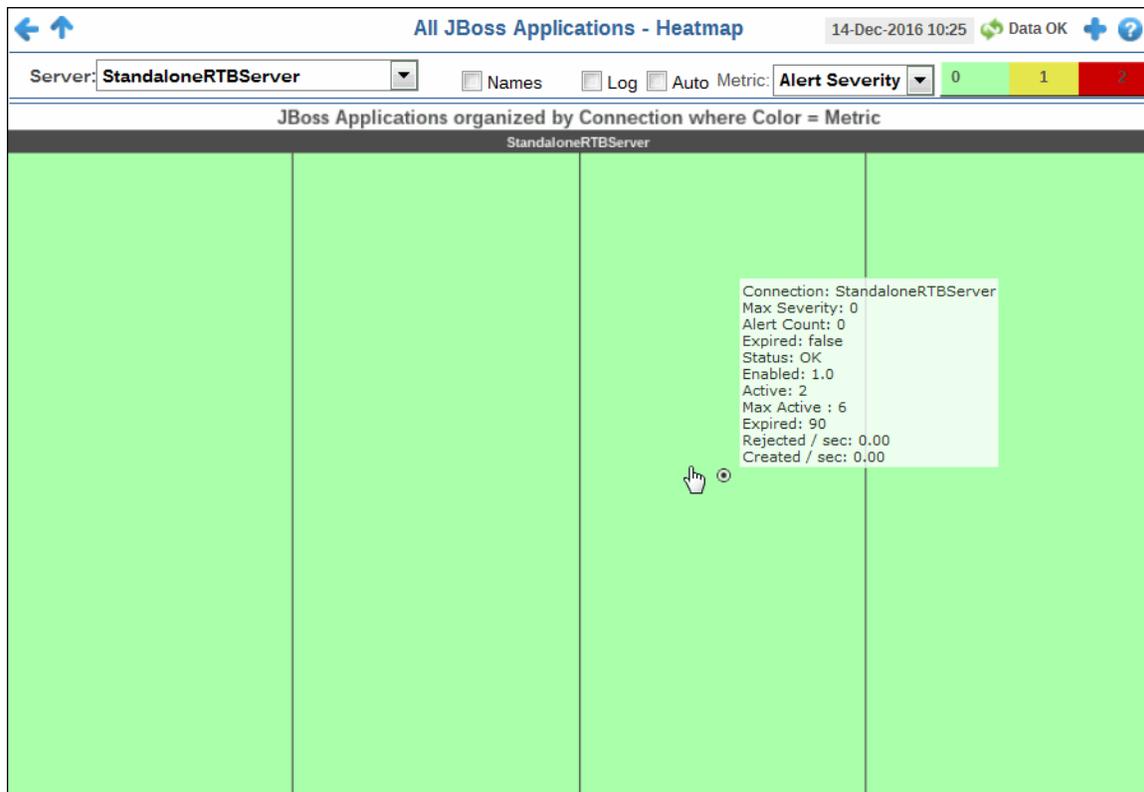
Displays in this View are:

- ["Applications Heatmap"](#)
- ["Applications Summary"](#)

**Applications Heatmap**

This heatmap shows the current status of all JBoss application connections. Use this to quickly identify the current session metrics for connections on one server or **All Servers**. Select a **Metric** from the drop-down menu, such as **Active Sessions** and **Average Alive Time**. By default, this display shows the **Alert Severity** metric.

Use the **Names** check-box  to include or exclude labels in the heatmap, or mouse over a rectangle to see additional metrics for connections. Clicking one of the rectangles in the heatmap opens the “Applications Summary” display, which allows you to see additional details for the selected server.



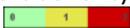
#### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- 6,047** The number of items currently in the display.

- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

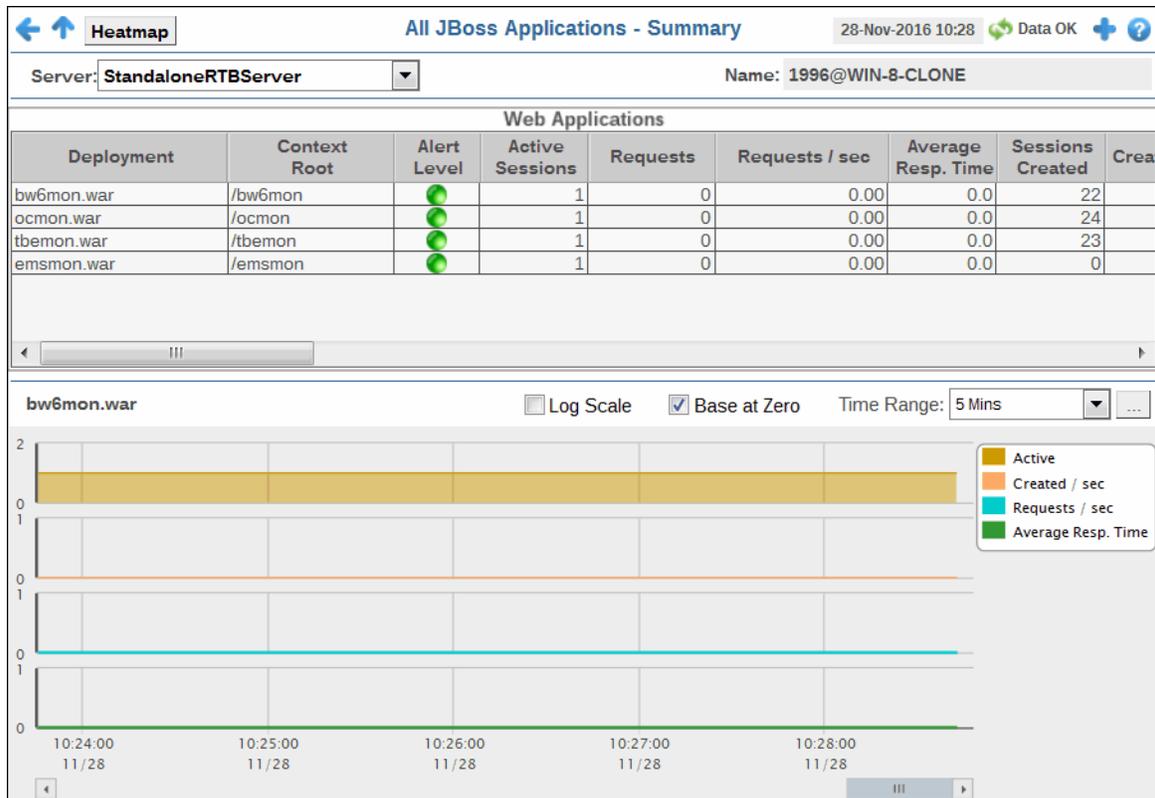
#### Fields and Data:

- Server** Choose a server to display.
- Names** Select this to include labels in the heatmap.
- Log** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Auto** Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value.  
**Note:** Some metrics auto-scale automatically, even when **Auto** is not selected.

<b>Metric</b>	Choose a metric to view in the display.
<b>Alert Severity</b>	<p>The current alert severity. Values range from <b>0</b> - <b>2</b>, as indicated in the color gradient  bar, where <b>2</b> is the highest Alert Severity:</p> <ul style="list-style-type: none"> <li><span style="color: red;">●</span> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li><span style="color: yellow;">●</span> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li><span style="color: green;">●</span> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Alert Count</b>	<p>The total number of critical and warning unacknowledged alerts. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.</p>
<b>Active Sessions</b>	<p>The number of currently active sessions. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined <b>JbossDeploymentActiveSessionsHigh</b> alert threshold. The middle value in the gradient bar indicates the middle value of the range.</p> <p>When <b>Auto</b> is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>
<b>Max Active Count</b>	<p>The total amount of active sessions. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum amount. The middle value in the gradient bar indicates the average amount.</p> <p>The <b>Auto</b> flag does not have any impact on this metric.</p>
<b>Avg Alive Time</b>	<p>The average amount of time a session lasts, in seconds. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum amount recorded. The middle value in the gradient bar indicates the average amount.</p>
<b>Created/sec</b>	<p>The average number of sessions created per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum number recorded. The middle value in the gradient bar indicates the average number.</p>
<b>Rejected/sec</b>	<p>The average number of rejected sessions per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum number recorded. The middle value in the gradient bar indicates the average number.</p>

## Applications Summary

Track current and historical performance of web applications on one server. Select a server from the **Server:** drop-down menu. Each row in the table is a different application on the selected server. Click a table row to populate the trend graphs.



### Title Bar (possible features are):

- ← ↑ Open the previous and upper display.
- ⊕ Open an instance of this display in a new window.
- ⓘ Open the online help page for this display.

Menu ▾, Table open commonly accessed displays.

6,047 The number of items currently in the display.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

Open the Alert Views - RTView Alerts Table display.

### Fields and Data

This display includes:

**Server** Choose a server to display.

**Name** The name of the connection.

### Web Application Table

Each table row is a different web application on the selected server.

**Deployment** The name of the **.war** file for the application.

<b>Context Root</b>	The location of the <b>.war</b> file for the application.
<b>Alert Level</b>	The current alert severity.  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
<b>Active Sessions</b>	The number of currently active sessions for the application.
<b>Requests/sec</b>	The number of requests per second for the application.
<b>Average Resp. Time</b>	The average response time for the application, in seconds.
<b>Sessions Created</b>	The total number of sessions created for the application.
<b>Created/sec</b>	The number of sessions created per second for the application.
<b>Sessions Rejected</b>	The total number of sessions rejected for the application.
<b>Rejected/sec</b>	The number of sessions rejected per second for the application.
<b>Max Active Sessions</b>	The maximum number of simultaneously active sessions counted for the application.
<b>Expired</b>	When checked, the monitor has not received monitoring data for the application in the defined time interval.
<b>Enabled</b>	Indicates whether the application has been enabled for use in the JBoss Server.
<b>Status</b>	Indicates the application status.
<b>Avg. Alive Time</b>	The average amount of time, in seconds,
<b>Max Alive Time</b>	The average amount of time, in seconds,
<b>Expired Sessions</b>	The number of expired sessions for the application.
<b>content</b>	The application content.
<b>Duplicated Session Ids</b>	The number of sessions containing duplicated session IDs.
<b>name</b>	The name of the <b>.war</b> file for the application.
<b>persistent</b>	Indicates whether
<b>Runtime Name</b>	The name of the <b>.war</b> file for the application.
<b>subsystem</b>	The subsystem of the application.
<b>Virtual Host</b>	The name of the virtual machine that hosts the application.
<b>Time Stamp</b>	The date and time of the last data update.
<b>Connection</b>	The connection name.

**Trend Graph**

Traces metrics for the selected server.

- Log Scale** Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Base at Zero** Use zero as the Y axis minimum for all graph traces.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar  .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

- Active** The number of currently active sessions for the application.
- Created/sec** The number of sessions created per second for the application.
- Requests/sec** The number of requests per second for the application.
- Average Resp. Time** The average response time for the application, in seconds.

---

## JBoss - HTML

This section describes the HTML version of the Solution Package for Red Hat JBoss.

The following Solution Package for Red Hat JBoss HTML version features an overview display, "[JBoss Overview - HTML](#)" (shown below), and the following Views which can be found under **Components** tab > **Application/Web Servers** > **JBoss**. For additional details, see vendor documentation.

- ["JBoss Servers - HTML"](#)
- ["JBoss Applications - HTML"](#)

## JBoss Overview - HTML

The JBoss Overview is the top-level display for the JBoss Solution Package, which provides a good starting point for immediately getting the status of all connections on your Data Server. You can select the RTView DataServer for which you want to see data and easily view the current data for that DataServer including:

- The total number of active alerts for the selected Data Server, including the total number of critical and warning alerts.
- The **Top 10 Process CPU%** hogs or **Threads** on your connected DataServer.

You can hover over each metric card to see more detail. You can also drill down to see even more detail by clicking on each respective region in the Overview.

The trend graph traces **Process CPU %**, **Active Sessions**, and **Sessions Created per second** and **Requests per second**. You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.

## JBoss Servers - HTML

Displays in this View are:

- ["JBoss Servers Table - HTML"](#)
- ["JBoss Servers Heatmap - HTML"](#)
- ["JBoss Single Server Summary - HTML"](#)

### JBoss Servers Table - HTML

Investigate detailed utilization metrics and configuration settings of JBoss servers. The **JBoss Servers Table** contains all metrics available for servers, including the number of current client connections.

Each row in the table contains data for a particular server. You can search, filter, sort and choose columns to include by clicking a column header icon (to the right of each column label) and selecting **Filter**, **Sort Ascending**, **Sort Descending** or **Columns**. Or just click a column header to sort.

Right-click on a table cell to **Export to Excel** or **Copy Cell Value**.

You can click on a row to drill-down to the ["JBoss Single Server Summary - HTML"](#) display and view details for that server.

### JBoss Servers Heatmap - HTML

View performance metrics for all monitored JBoss servers. The heatmap organizes JBoss servers by connection, and uses color to show the most critical Metric value for each JBoss connection. Each rectangle in the heatmap represents a connection.

Use this display to see at-a-glance the health of all your JBoss servers. You can select the heatmap color metric from a list including **Process CPU Load %**, **System CPU Load %** and **Memory Used %**.

Use the check-boxes  to include or exclude labels in the heatmap. Move your mouse over a rectangle to see additional information. Drill-down and investigate by clicking a rectangle in the heatmap to view details for the selected Web module in the ["JBoss Single Server Summary - HTML"](#) display.

## JBoss Single Server Summary - HTML

Track utilization and performance metrics for one JBoss server. Clicking on the metric boxes at the top of the display takes you to the ["JBoss Servers Table - HTML"](#) display, where you can compare and sort performance values against other JBoss servers.

The trend graph traces **Process CPU %**, **Active Sessions**, **Sessions Created per second** and **Requests per second**. You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.

Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display.

## JBoss Applications - HTML

Displays in this View are:

- ["JBoss Apps Table - HTML"](#)
- ["JBoss Apps Heatmap - HTML"](#)
- ["JBoss App Totals - HTML"](#)

## JBoss Apps Table - HTML

Investigate detailed utilization metrics for all JBoss applications. This display contains all metrics available for JBoss applications.

Each row in the table contains data for a particular web module. You can search, filter, sort and choose columns to include by clicking a column header icon (to the right of each column label) and selecting **Filter**, **Sort Ascending**, **Sort Descending** or **Columns**.

Or just click a column header to sort. Right-click on a table cell to **Export to Excel** or **Copy Cell Value**.

To investigate further, double-click a web module to see details in the ["JBoss App Totals - HTML"](#) display.

## JBoss Apps Heatmap - HTML

This heatmap allows you to view the status and alerts of JBoss applications.

Use the **Metric** drop-down menu to view the **Alert Severity**, **Alert Count**, **Max Active Count**, **Average Alive Time**, **Created per Second** or **Rejected per Second**.

Click on a rectangle to drill-down to the ["JBoss App Totals - HTML"](#) display and view metrics for a particular web module. Toggle between the commonly accessed Table and Heatmap displays by clicking the drop down list on the display title.

Mouse-over rectangles to view more details about host performance and status.

You can view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.

## JBoss App Totals - HTML

Track utilization and performance metrics for a particular JBoss application. Clicking on the sessions/processing rate information boxes at the top of the display takes you to the "[JBoss Apps Table - HTML](#)" display, where you can compare and sort performance values against other JBoss applications.

Use the table to compare detailed utilization metrics for all web applications. Each row in the table contains data for a particular web application. You can search, filter, sort and choose columns to include by clicking a column header icon (to the right of each column label) and selecting **Filter**, **Sort Ascending**, **Sort Descending** or **Columns**.

Or just click a column header to sort. Right-click on a table cell to **Export to Excel** or **Copy Cell Value**.

The trend graph traces **Active Sessions**, **Sessions Created** and **Requests per second**. You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.

Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display.

---

## MongoDB

The following MongoDB Views (and their associated displays) can be found under **Components** tab > **Databases** > **MongoDB**:

- "[Mongo Instance View](#)"
- "[Mongo Database View](#)"
- "[Mongo Collection View](#)"

## Mongo Instance View

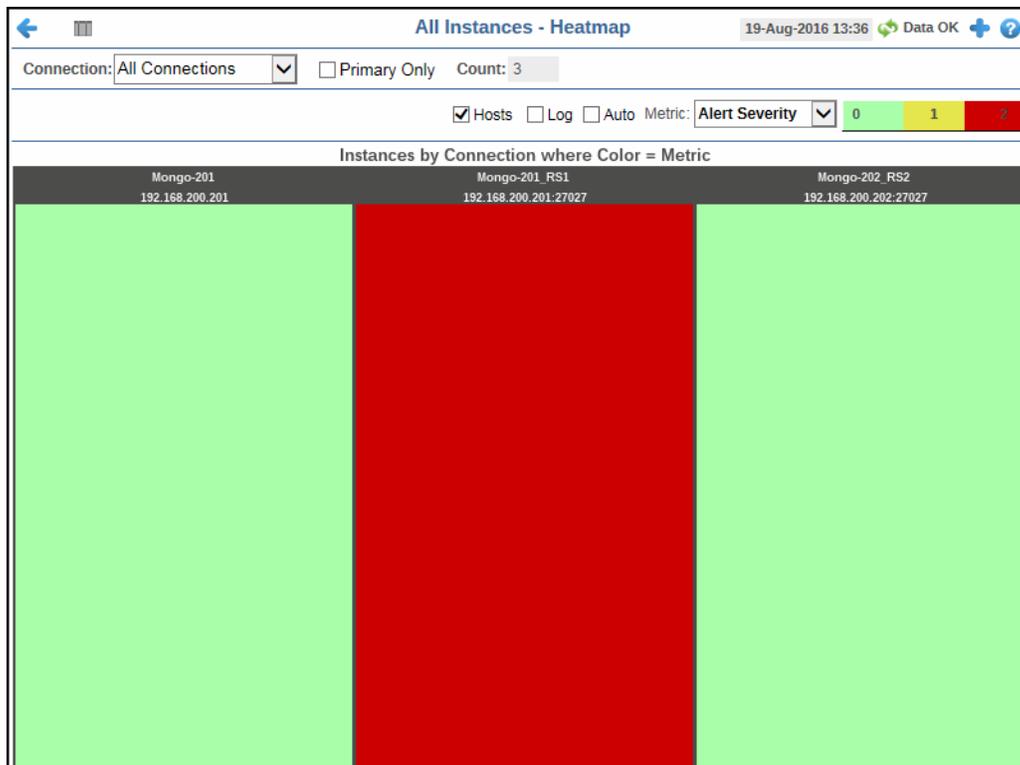
These displays present performance metrics and alert statuses for all MongoDB instances. The first two displays show different views of the same data:

- "[All Instances Heatmap](#)": This heatmap shows status and alerts for all MongoDB instances.
- "[All Instances Table](#)": This table shows all available utilization metrics for all MongoDB instances.
- "[Single Instance Summary](#)": This summary enables you to view available utilization metrics for a single MongoDB instance.

## All Instances Heatmap

View status and alerts of all MongoDB Instances. Use the **Metric** drop-down menu to view the **Alert Severity**, **Alert Count**, **Physical Memory**, **Open Cursors**, **Connections**, or **Databases**.

The heatmap is organized by host, each rectangle representing a connection. The rectangle color indicates the most critical alert state. Click on a node to drill-down to the “[Single Instance Summary](#)” display and view metrics for a particular connection. You can toggle between the commonly accessed **Table** and **Heatmap** displays by clicking the icon in the upper left-hand corner. Mouse-over rectangles to view more details about host performance and status.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- 6,047** The number of items currently in the display.

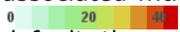
- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

### Fields and Data

This display includes:

- Connection** Select the connection from the drop down list for which you want to view data.

- Primary Only** Selecting this check box displays connections in the heatmap that have **Designation in Set** (within a replica set) defined as **Primary**, as well as those connections that are not part of a replica set (do not have a defined **Designation in Set**). Those connections with **Designation in Set** defined as **Secondary** will not be displayed.
- Count** The total number of active, inactive, and standby connections.
- Hosts** Select this check box to display the IP address of the host for each rectangle.
- Log** This option enables visualization on a logarithmic scale, and should be used when the range in your data is very broad. For example, if you have data that ranges from the tens to the thousands, then data in the range of tens will be neglected visually if you do not check this option. This option makes data on both extreme ranges visible by using the logarithmic of the values rather than the actual values.
- Auto** When checked, the values of the selected metric are auto-scaled to its highest defined value. When unchecked, the values of the selected metric display based on the threshold defined for the alert associated with the selected metric. Selecting Auto helps to visualize the range of the values currently present for the selected metric instead of the threshold of the alert that has been associated with the metric. All metrics that have not been associated in the heatmap defaults with alerts use a monochromatic color gradient bar (whites and greens). All metrics that have been associated in the heatmap defaults with alerts use a multi-chromatic color gradient bar (reds, yellows, white, and greens).
- Metric** Select the metric driving the heatmap display. The default is Alert Severity. Each **Metric** has a color gradient bar that maps values to colors. The heatmap organizes the instances by host, where each rectangle represents an instance. Mouse-over any rectangle to display the current values of the metrics for the instance. Click on a rectangle to drill-down to the associated ["Single Instance Summary"](#) display for a detailed view of metrics for that particular instance.
- Alert Severity** The maximum alert level in the item (index) associated with the rectangle. Values range from **0** to **2**, as indicated in the color gradient bar , where **2** is the greatest **Alert Severity**.
- 2** -- Metrics that have exceeded their specified **ALARMLEVEL** threshold and have an Alert Severity value of **2** are shown in red. For a given rectangle, this indicates that one or more metrics have exceeded their alarm threshold.
- 1** -- Metrics that have exceeded their specified **WARNINGLEVEL** threshold and have an Alert Severity value of **1** are shown in yellow. For a given rectangle, this indicates that one or more metrics have exceeded their warning threshold.
- 0** -- Metrics that have not exceeded either specified threshold have an Alert Severity value of **0** and are shown in green. For a given rectangle, this indicates that no metrics have exceeded a specified alert threshold.
- Alert Count** The total number of alarm and warning alerts in a given item (index) associated with the rectangle.
- The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
- Physical Memory** The total amount of physical memory currently being used in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum amount of physical memory in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
- The **Auto** option does not impact this metric.

<b>Open Cursors</b>	<p>The total number of open cursors in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from <b>0</b> to the alert threshold of <b>MongoInstanceOpenCursorsHigh</b>, which is <b>2000</b>. The middle value in the gradient bar indicates the middle value of the range (the default is <b>1000</b>).</p> <p>When <b>Auto</b> is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>
<b>Connections</b>	<p>The total number of connections in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of connections in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p> <p>The <b>Auto</b> option does not impact this metric.</p>
<b>Databases</b>	<p>The total number of databases in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of databases in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p> <p>The <b>Auto</b> option does not impact this metric.</p>

## All Instances Table

This display enables you to investigate detailed utilization metrics for all MongoDB Instances. The **All Instances Table** contains all metrics available for instances, including the number of current connections. Each row in the table contains data for a particular connection. Click a column header to sort column data in numerical or alphabetical order. Click on a table row to drill-down to the ["Single Instance Summary"](#) display and view metrics for that particular instance. You can click the icon in the upper left-hand corner to toggle between the commonly accessed **Table** and **Heatmap** displays.

Connection	Host	Conn Status	Alert Level	Alert Count	Average Back Flush	Connections	Databases	Designation In Set
Mongo-201	192.168.200.201			0	0.00	27	2	

**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- The number of items currently in the display.
- Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

**Fields and Data**

This display includes:

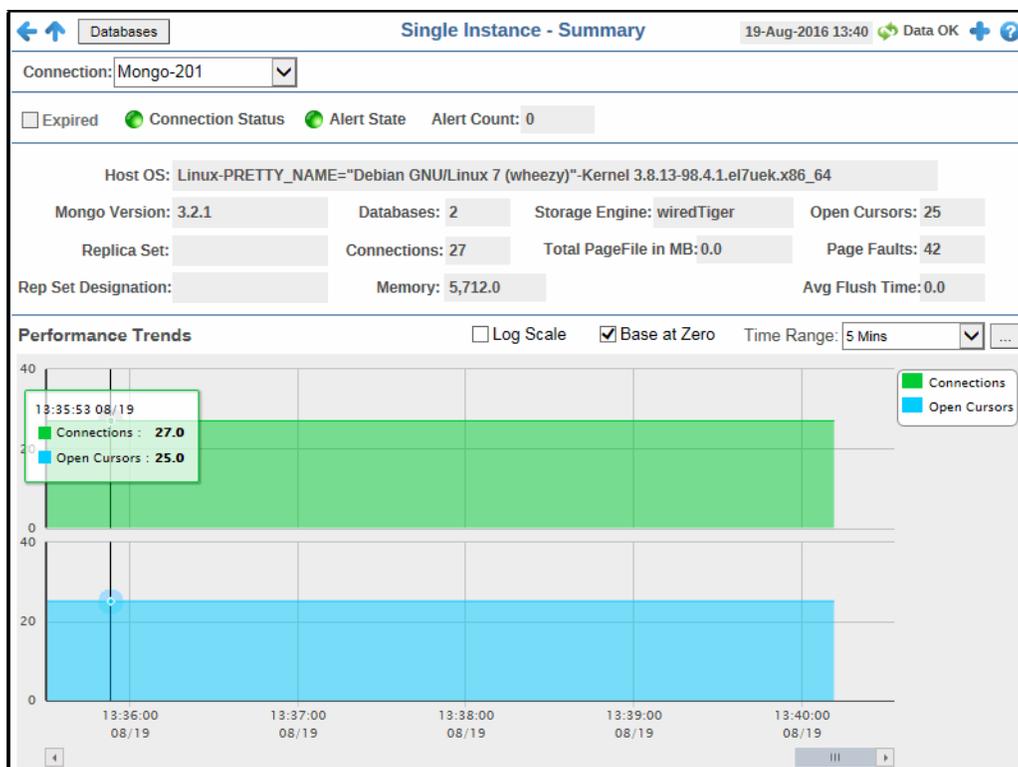
- Connection** Select the connection for which you want to view data, or select **All Connections** to view data for all connections.
- Primary Only** Selecting this check box displays connections in the table that have **Designation in Set** (within a replica set) defined as **Primary**, as well as those connections that are not part of a replica set (do not have a defined **Designation in Set**). Those connections with **Designation in Set** defined as **Secondary** will not be displayed.
- Count** The total number of connections displayed in the table.
- Table** This table shows information for the selected connection(s). Click on a table row to drill-down to the ["Single Instance Summary"](#) display and view metrics for that particular server.
  - Connection** The name of the connection.

<b>Host</b>	The host name returned by MongoDB or the host provided by the user to use for connection if the host is not available.
<b>Conn Status</b>	The connection status of the Connection/Host.  -- The host is not connected.  -- The host is partially connected, which occurs when the connection has succeeded but the credentials given do not allow access to certain metrics.  -- The host is connected.
<b>Alert Level</b>	The current alert level.  -- One or more alerts have exceeded their specified <b>ALARMLEVEL</b> threshold.  -- One or more alerts have exceeded their specified <b>WARNINGLEVEL</b> threshold.  -- No alerts have exceeded an alert threshold.
<b>Alert Count</b>	The total number of alerts for the connection.
<b>Average Back Flush</b>	The average time, in milliseconds, for each flush to disk, calculated by dividing the total number of milliseconds by the total number of flushes.  <b>Note:</b> Background flushing information only appears for instances using the This metric only displays when the storage engine is <b>MMAPv1</b> storage engine.
<b>Connections</b>	The number of connections coming in from the clients to the database server, including the current monitor session.
<b>Databases</b>	The number of databases being hosted by the instance.
<b>Designation In Set</b>	The designation of this member of the replica set ( <b>primary/secondary</b> ). This column will be empty if no replica set is configured, or set to <b>unknown</b> if there is no connection.
<b>Host OS Version</b>	The version of the operating system used by the host.
<b>How Long As Primary</b>	The amount of time the instance has been a primary instance. This field is only populated for primary instances.
<b>MongoDB Version</b>	The version number of the mongod instance.
<b>Open Cursors</b>	The total number of open cursors for the connection.
<b>Ops Log Lag</b>	The amount of time (in hours:minutes:seconds) in which the secondary instance is behind the primary instance. This field is only populated for secondary instances.
<b>Ops Log Length</b>	The length of the OpsLog collection, in bytes.
<b>Page Faults</b>	The number of page faults for the connection. MongoDB reports its triggered page faults as the total number of page faults in one second.
<b>Physical Memory MB</b>	The total amount of system memory (RAM), in megabytes.
<b>ReplicaSet</b>	The name of the replica set in which the mongod is a part of, if configured. This column will be empty if no replica set is configured, or set to <b>unknown</b> if there is no connection. All hosts in the replica set must have the same set name.

<b>Storage Engine</b>	The name of the current storage engine. The name can be either <b>MMAPv1</b> or <b>WiredTiger</b> . <b>WiredTiger</b> is the default as of MongoDB version 3.2.
<b>Total Page File MB</b>	The total size of pagefile defined for the connection, in megabytes. This metric only displays when the storage engine is <b>MMAPv1</b> .
<b>Uptime</b>	The amount of time since the instance was last started, shown in days, hours, and minutes (for example, 1d 23:43).
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > ( <b>Project Name</b> ) > <b>Solution Package Configuration</b> > <b>MongoDB</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Time Stamp</b>	The date and time this row of data was last updated.

## Single Instance Summary

Track utilization and performance metrics for specific instances.



**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- The number of items currently in the display.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

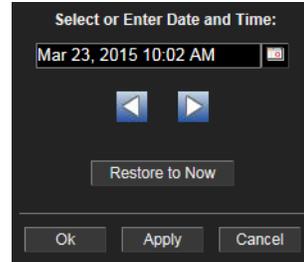
Open the **Alert Views - RTView Alerts Table** display.

## Fields and Data

This display includes:

<b>Connection</b>	Select the connection for which you want to view data.
<b>Expired</b>	This check box becomes automatically checked when the data has exceeded the specified cache expiration time (set by default at 45 seconds) and is no longer current. Once the cache has been refreshed and is displaying current data, the check box will return to being unchecked. This check box will remain unchecked as long as the cache has been refreshed within the specified cache expiration time and the data is current.
<b>Connection Status</b>	The connection status of the Connection/Host. <ul style="list-style-type: none"> <li> -- The host is not connected.</li> <li> -- The host is partially connected, which occurs when the connection has succeeded but the credentials given do not allow access to certain metrics.</li> <li> -- The host is connected.</li> </ul>
<b>Alert State</b>	The current alert level. <ul style="list-style-type: none"> <li> -- One or more alerts have exceeded their specified <b>ALARMLEVEL</b> threshold.</li> <li> -- One or more alerts have exceeded their specified <b>WARNINGLEVEL</b> threshold.</li> <li> -- No alerts have exceeded an alert threshold.</li> </ul>
<b>Alert Count</b>	The total number of alerts for the connection.
<b>Host OS</b>	The version of the operating system used by the host.
<b>Mongo Version</b>	The version number of the mongod instance.
<b>Databases</b>	The number of databases being hosted by the instance.
<b>Storage Engine</b>	The name of the current storage engine. The name can be either <b>MMAPv1</b> or <b>WiredTiger</b> . <b>WiredTiger</b> is the default as of MongoDB version 3.2.
<b>Open Cursors</b>	The total number of open cursors for the connection.
<b>Replica Set</b>	The name of the replica set in which the mongod is a part of, if configured. This column will be empty if no replica set is configured, or set to <b>unknown</b> if there is no connection. All hosts in the replica set must have the same set name.
<b>Connections</b>	The number of connections coming in from the clients to the database server, including the current monitor session.
<b>Total PageFile in MB</b>	The total size of pagefile defined for the connection, in megabytes. This metric only displays when the storage engine is <b>MMAPv1</b> .
<b>Page Faults</b>	The number of page faults for the connection. MongoDB reports its triggered page faults as the total number of page faults in one second.

<b>Rep Set Designation</b>	The designation of this member of the replica set ( <b>primary/secondary</b> ). This column will be empty if no replica set is configured, or set to <b>unknown</b> if there is no connection.
<b>Memory</b>	The total amount of system memory (RAM), in megabytes.
<b>Avg Flush Time</b>	The average time, in milliseconds, for each flush to disk, calculated by dividing the total number of milliseconds by the total number of flushes. <b>Note:</b> Background flushing information only appears for instances using the <b>MMAPv1</b> storage engine.
<b>Performance Trends Graph</b>	Shows connection and open cursor data for the connection. <b>Connections</b> -- Traces the total number of connections coming in from the clients. <b>Open Cursors</b> -- Traces the total number of open cursors on the connection.
<b>Log Scale</b>	This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.
<b>Base at Zero</b>	When this option is checked, zero is set as the Y axis minimum for all graph traces.
<b>Time Range</b>	Select a time range from the drop down menu varying from <b>2 Minutes</b> to <b>Last 7 Days</b> , or display <b>All Data</b> . To specify a time range, click the  button.



By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Mongo Database View

These displays present detailed performance metrics and alert statuses for all databases (in a heatmap or a tabular format) or for an individual database.

- ["All Databases Heatmap"](#): Displays a heatmap view of alert states for all databases.
- ["All Databases Table"](#): Displays a tabular view of all databases and their associated metrics for a single connection, or of all databases and their associated metrics for all connections.
- ["Database Summary"](#): Displays metrics for a specific database.

---

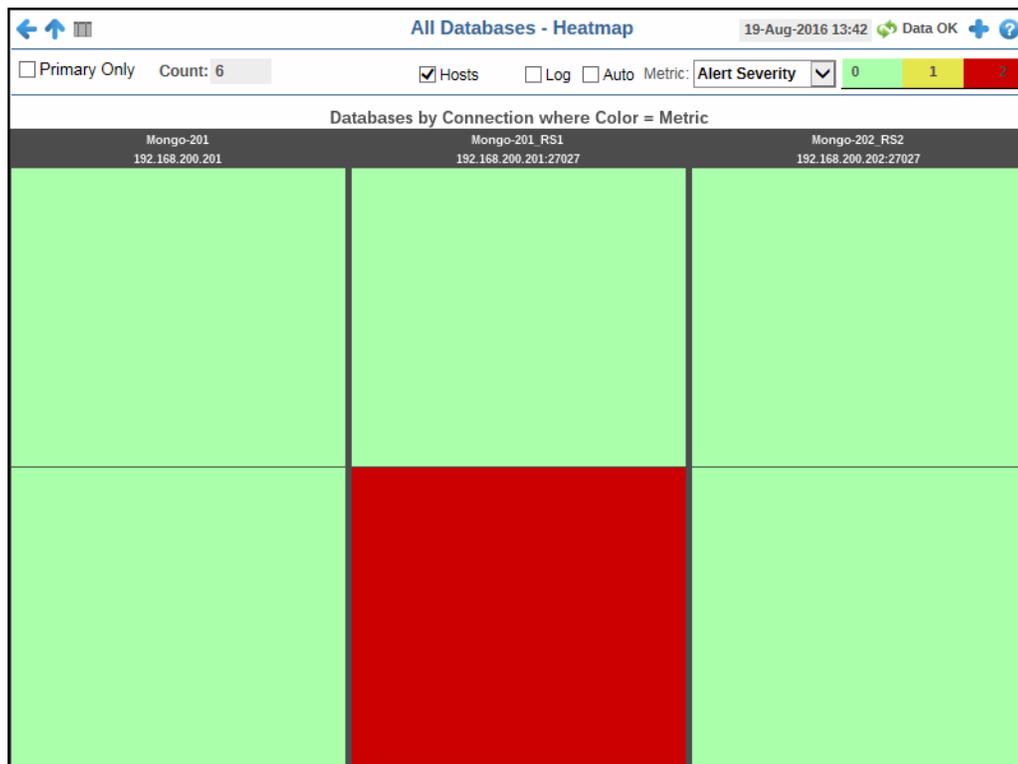
**Note:** No database information will display in the heatmap, table, or summary displays if a connection cannot be established.

---

### All Databases Heatmap

Track utilization and performance metrics for all databases in a heatmap format. Use the **Metric** drop-down menu to view **Alert Severity**, **Alert Count**, **NumObjects** (number of objects), or **AvgObjectSize** (average object size).

The heatmap is organized so that each rectangle represents a database associated with a specific connection. The rectangle color indicates the value of the selected metric in the **Metric** drop down list. You can mouse-over rectangles to view more details about the performance and status of each database or click on a rectangle to drill-down to the ["Database Summary"](#) display and view metrics for that particular database. You can click the table icon  in this display to navigate to the ["All Databases Table"](#) display.

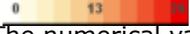
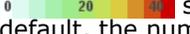




## Fields and Data

This display includes:

- Primary Only** Selecting this check box displays connections in the heatmap that have **Designation in Set** (within a replica set) defined as **Primary**, as well as those connections that are not part of a replica set (do not have a defined **Designation in Set**). Those connections with **Designation in Set** defined as **Secondary** will not be displayed.
- Count** The total number of active and inactive databases.
- Hosts** Select this check box to display the name/IP address of the host for each rectangle.
- Log** This option enables visualization on a logarithmic scale, and should be used when the range in your data is very broad. For example, if you have data that ranges from the tens to the thousands, then data in the range of tens will be neglected visually if you do not check this option. This option makes data on both extreme ranges visible by using the logarithmic of the values rather than the actual values.
- Auto** When checked, the values of the selected metric are auto-scaled to its highest defined value. When unchecked, the values of the selected metric display based on the threshold defined for the alert associated with the selected metric. Selecting **Auto** helps to visualize the range of the values currently present for the selected metric instead of the threshold of the alert that has been associated with the metric. All metrics that have not been associated in the heatmap defaults with alerts use a monochromatic color gradient bar (whites and greens). All metrics that have been associated in the heatmap defaults with alerts use a multi-chromatic color gradient bar (reds, yellows, white, and greens).
- Metric** Select the metric driving the heatmap display. The default is **Alert Severity**. Each **Metric** has a color gradient bar that maps values to colors. The heatmap organizes the databases by connection, where each rectangle represents a database. Mouse-over any rectangle to display the current values of the metrics for the database. Click on a rectangle to drill-down to the associated ["Database Summary"](#) display for a detailed view of metrics for that particular database.

- Alert Severity** The maximum alert level in the item (index) associated with the rectangle. Values range from **0** to **2**, as indicated in the color gradient bar , where **2** is the greatest **Alert Severity**.
- 2** -- Metrics that have exceeded their specified **ALARMLEVEL** threshold and have an Alert Severity value of **2** are shown in red. For a given rectangle, this indicates that one or more metrics have exceeded their alarm threshold.
- 1** -- Metrics that have exceeded their specified **WARNINGLEVEL** threshold and have an Alert Severity value of **1** are shown in yellow. For a given rectangle, this indicates that one or more metrics have exceeded their warning threshold.
- 0** -- Metrics that have not exceeded either specified threshold have an Alert Severity value of **0** and are shown in green. For a given rectangle, this indicates that no metrics have exceeded a specified alert threshold.
- Alert Count** The total number of alarm and warning alerts in a given item (index) associated with the rectangle.
- The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
- Collections** The total number of collections in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum number of collections in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
- The **Auto** option does not impact this metric.
- Data Size** The total size (in bytes) of the data in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of **MongoDatabaseDataSizeHigh**, which is **100,000**. The middle value in the gradient bar indicates the middle value of the range (the default is **50,000**).
- When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

## All Databases Table

View details for all databases in a single connection, or view details for all databases in all connections.

Connection	Host	Name	Alert Level	Alert Count	Status UpDown	Collections	Data Size	File Size
Mongo-201	192.168.200.201	local	🟢	0	☑	1	7,596	0
Mongo-201	192.168.200.201	meantest	🟢	0	☑	1	390	0

### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- The number of items currently in the display.

- Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

### Fields and Data

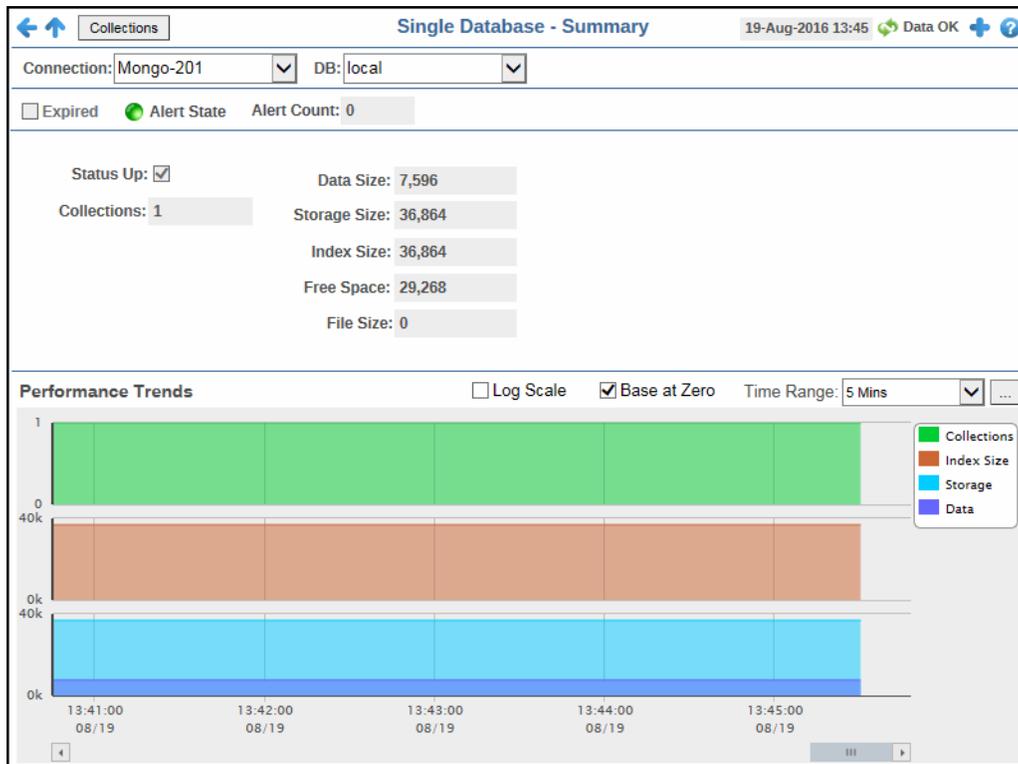
This display includes:

- Connection** Select the connection for which you want to view data, or select **All Connections** to view data for all connections.
- Primary Only** Selecting this check box displays connections in the table that have **Designation in Set** (within a replica set) defined as **Primary**, as well as those connections that are not part of a replica set (do not have a defined **Designation in Set**). Those connections with **Designation in Set** defined as **Secondary** will not be displayed.
- Count** The total number of databases displayed in the table.

<b>Table</b>	This table shows information for the selected connection(s). Click on a table row to drill-down to the <a href="#">"Database Summary"</a> display and view metrics for that particular server.
<b>Connection</b>	The name of the connection
<b>Host</b>	The host name returned by MongoDB or the host provided by the user to use for connection if the host is not available.
<b>Name</b>	The name of the database.
<b>Alert Level</b>	The current alert level.  -- One or more alerts have exceeded their specified <b>ALARMLEVEL</b> threshold.  -- One or more alerts have exceeded their specified <b>WARNINGLEVEL</b> threshold.  -- No alerts have exceeded an alert threshold.
<b>Alert Count</b>	The total number of alerts for the database.
<b>Status UpDown</b>	When checked, signifies that the database is up and running.
<b>Collections</b>	The number of collections in the database.
<b>Data Size</b>	The total size, in bytes, of the data held in the database including the padding factor. The <b>Data Size</b> will not decrease when the document size decreases, but will decrease when documents are removed. <b>Note:</b> The <b>scale</b> argument affects this value.
<b>File Size</b>	The total size, in bytes, of the data files in the database. This value includes preallocated space as well as the padding factor, and only reflects the size of the data files in the database and not the size of the namespace file.
<b>Free Space</b>	The total free space remaining on the database ( <b>Storage Size</b> minus <b>Data Size</b> ).
<b>Index Size</b>	The total size, in bytes, of all indexes created on the database. <b>Note:</b> The <b>scale</b> argument affects this value.
<b>Storage Size</b>	The total amount of space, in bytes, allocated to collections in this database for document storage. The Storage Size does not decrease when documents are removed or the size of the documents decrease. <b>Note:</b> The <b>scale</b> argument affects this value.
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > <b>(Project Name)</b> > <b>Solution Package Configuration</b> > <b>MongoDB</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>time_stamp</b>	The date and time the data in this row was last updated.

## Database Summary

View all available utilization and performance data for a specific database.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

### Fields and Data

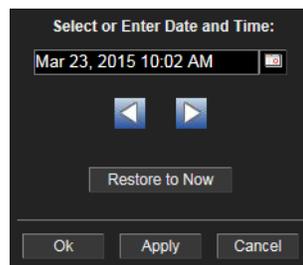
This display includes:

- Connection** Select the connection for which you want to view data.
- DB** Select the database for which you want to view data.
- Expired** When checked, performance data has not been received within the time specified (in seconds) in the **Expire Time** field in the **Duration** region in the RTView Configuration Application > **(Project Name)** > **Solution Package Configuration** > **MongoDB** > **DATA STORAGE** tab. The **Delete Time** field (also in the **Duration** region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.

<b>Alert State</b>	<p>The current alert level.</p> <ul style="list-style-type: none"> <li>● -- One or more alerts have exceeded their specified <b>ALARMLEVEL</b> threshold.</li> <li>● -- One or more alerts have exceeded their specified <b>WARNINGLEVEL</b> threshold.</li> <li>● -- No alerts have exceeded an alert threshold.</li> </ul>
<b>Alert Count</b>	The total number of alerts for the database.
<b>Status Up</b>	When checked, signifies that the database is up and running.
<b>Collections</b>	The total number of collections in the database.
<b>Data Size</b>	<p>The total size, in bytes, of the data held in the database including the padding factor. The <b>Data Size</b> will not decrease when the document size decreases, but will decrease when documents are removed.</p> <p><b>Note:</b> The <b>scale</b> argument affects this value.</p>
<b>Storage Size</b>	<p>The total amount of space, in bytes, allocated to collections in this database for document storage. The Storage Size does not decrease when documents are removed or the size of the documents decrease.</p> <p><b>Note:</b> The <b>scale</b> argument affects this value.</p>
<b>Index Size</b>	<p>The total size, in bytes, of all indexes created on the database.</p> <p><b>Note:</b> The <b>scale</b> argument affects this value.</p>
<b>Free Space</b>	The total free space remaining on the database ( <b>Storage Size</b> minus <b>Data Size</b> ).
<b>File Size</b>	The total size, in bytes, of the data files in the database. This value includes preallocated space as well as the padding factor, and only reflects the size of the data files in the database and not the size of the <b>namespace</b> file.
<b>Performance Trends Graph</b>	<p>Shows connection and open cursor data for the connection.</p> <ul style="list-style-type: none"> <li><b>Collections</b> -- Traces the total number of collections in the database.</li> <li><b>Index Size</b> -- Traces the total size of indexes created on the database.</li> <li><b>Storage</b> -- Traces the total amount of space allocated to collections in the database.</li> <li><b>Data</b> -- Traces the total size of the data held in the database.</li> </ul> <p><b>Log Scale</b> This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.</p>

**Base at Zero** When this option is checked, zero is set as the Y axis minimum for all graph traces.

**Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows  to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Mongo Collection View

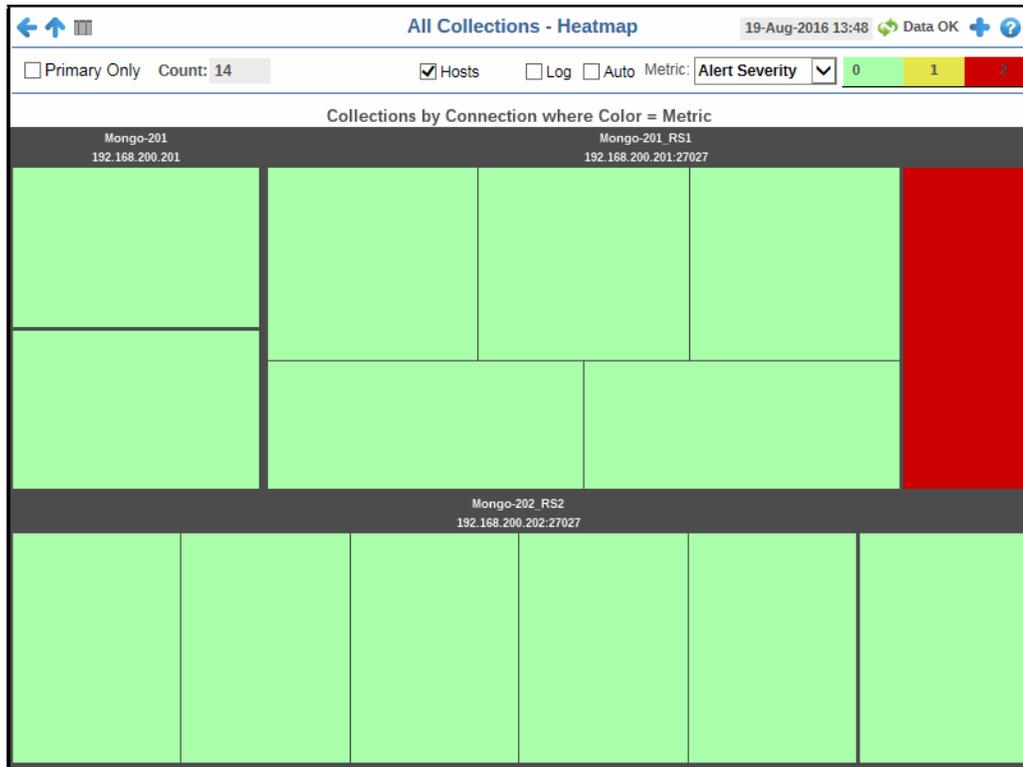
These displays present several views of performance metrics for collections. You can view heatmap or tabular views of all collections that exist in the connections in the [“All Collections Heatmap”](#) and [“All Collections Table”](#) displays, or you can view all details for a specific collection contained in a particular database in the [“Collection Summary”](#) display.

- [“All Collections Heatmap”](#): A heatmap representation that allows you to view performance and utilization metrics for all collections that exist in each of your connections.
- [“All Collections Table”](#): A tabular view that allows you to view performance and utilization metrics for all collections in a particular database, or for all collections on all databases.
- [“Collection Summary”](#): Shows detailed performance and utilization metrics and trends for a specified collection on a particular database.

### All Collections Heatmap

This display provides a heatmap view of the status and alerts of all collections within each connection. Use the **Metric** drop-down menu to view **Alert Severity**, **Alert Count**, **NumObjects** (number of objects), or **AvgObjectSize** (average object size).

The heatmap is organized so that each rectangle represents a collection contained within a specific connection. The rectangle color indicates the value of the selected metric in the **Metric** drop down list. You can mouse-over rectangles to view more details about the performance and status of each collection or click on a rectangle to drill-down to the [“Collection Summary”](#) display and view metrics for that particular collection. You can click the table icon  in this display to navigate to the [“All Collections Table”](#) display.



**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- The number of items currently in the display.

- Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

**Fields and Data**

This display includes:

- Primary Only** Selecting this check box displays connections in the heatmap that have **Designation in Set** (within a replica set) defined as **Primary**, as well as those connections that are not part of a replica set (do not have a defined **Designation in Set**). Those connections with **Designation in Set** defined as **Secondary** will not be displayed.
- Count** The total number of collections.
- Hosts** Select this check box to display the names of the hosts in the heatmap.

**Log** This option enables visualization on a logarithmic scale, and should be used when the range in your data is very broad. For example, if you have data that ranges from the tens to the thousands, then data in the range of tens will be neglected visually if you do not check this option. This option makes data on both extreme ranges visible by using the logarithmic of the values rather than the actual values.

**Auto** When checked, the values of the selected metric are auto-scaled to its highest defined value. When unchecked, the values of the selected metric display based on the threshold defined for the alert associated with the selected metric. Selecting **Auto** helps to visualize the range of the values currently present for the selected metric instead of the threshold of the alert that has been associated with the metric. All metrics that have not been associated in the heatmap defaults with alerts use a monochromatic color gradient bar (whites and greens). All metrics that have been associated in the heatmap defaults with alerts use a multi-chromatic color gradient bar (reds, yellows, white, and greens).

**Metric** Select the metric driving the heatmap display. The default is **Alert Severity**. Each **Metric** has a color gradient bar that maps values to colors. The heatmap organizes the collections by connection, where each rectangle represents a collection. Mouse-over any rectangle to display the current values of the metrics for the collection. Click on a rectangle to drill-down to the associated "[Collection Summary](#)" display for a detailed view of metrics for that particular collection.

**Alert Severity**

The maximum alert level in the item (index) associated with the rectangle. Values range from **0** to **2**, as indicated in the color gradient bar , where **2** is the greatest **Alert Severity**.

**2** -- Metrics that have exceeded their specified **ALARMLEVEL** threshold and have an Alert Severity value of **2** are shown in red. For a given rectangle, this indicates that one or more metrics have exceeded their alarm threshold.

**1** -- Metrics that have exceeded their specified **WARNINGLEVEL** threshold and have an Alert Severity value of **1** are shown in yellow. For a given rectangle, this indicates that one or more metrics have exceeded their warning threshold.

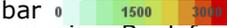
**0** -- Metrics that have not exceeded either specified threshold have an Alert Severity value of **0** and are shown in green. For a given rectangle, this indicates that no metrics have exceeded a specified alert threshold.

**Alert Count**

The total number of alarm and warning alerts in a given item (index) associated with the rectangle.

The color gradient bar , shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

**NumObjects**

The total number of objects or documents in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of **MongoCollectionNumObjectsHigh**, which is **2000**. The middle value in the gradient bar indicates the middle value of the range (the default is **1000**).

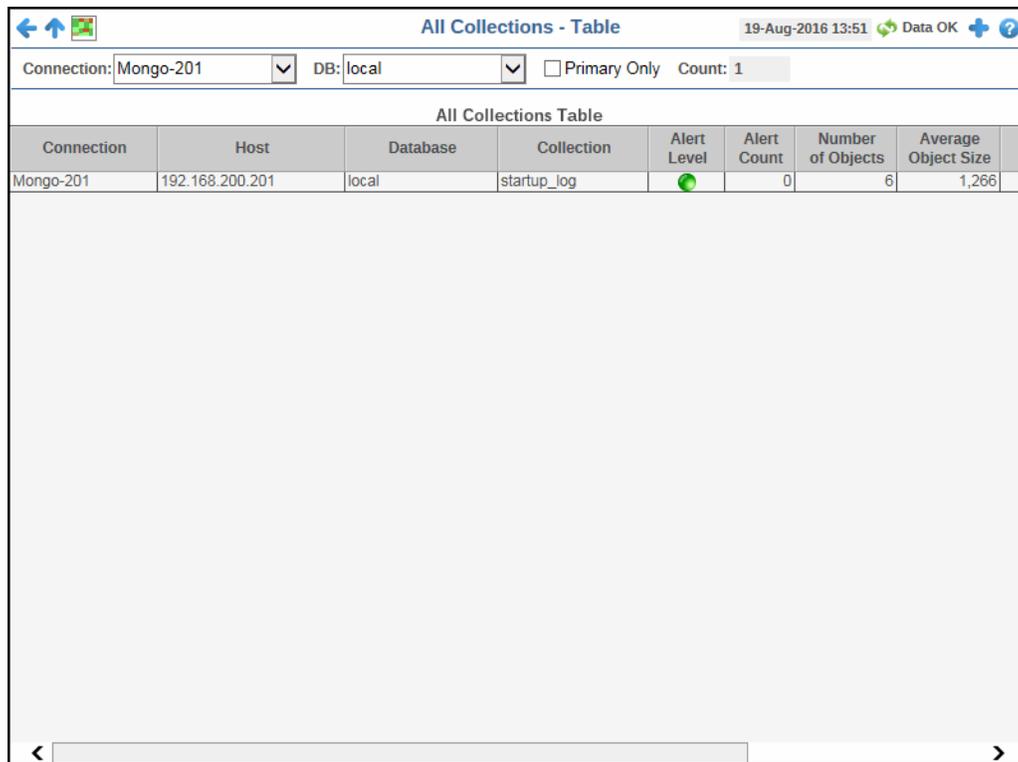
When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

**AvgObjectSize**

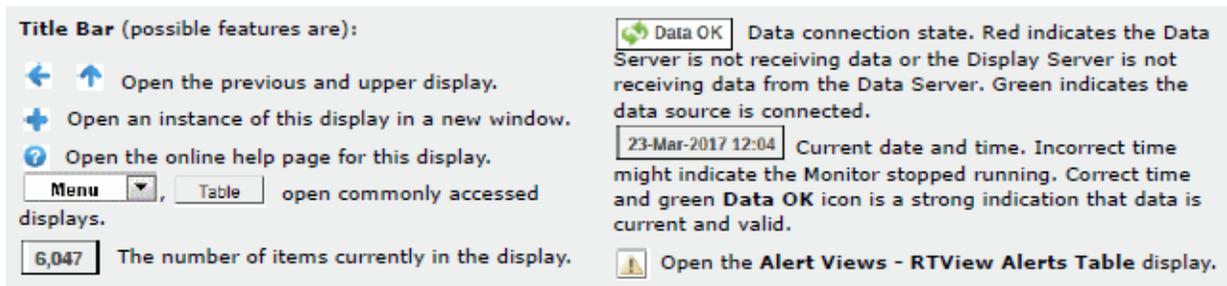
The average size (in bytes) of an object in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of objects in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

**All Collections Table**

Track performance and utilization metrics for all collections on a single database, or for all connections on all databases.



All Collections Table							
Connection	Host	Database	Collection	Alert Level	Alert Count	Number of Objects	Average Object Size
Mongo-201	192.168.200.201	local	startup_log	0	0	6	1,266



## Fields and Data

This display includes:

- Connection** Select the connection for which you want to view collection data.
- DB** Select the database for which you want to view collection data, or select **All Databases** to view all collections for all databases.
- Primary Only** Selecting this check box displays connections in the table that have **Designation in Set** (within a replica set) defined as **Primary**, as well as those connections that are not part of a replica set (do not have a defined **Designation in Set**). Those connections with **Designation in Set** defined as **Secondary** will not be displayed.
- Count** The total number of collections found for the selected database(s).
- All Collections Table** This table describes all topics on the selected server. Click a row to view metrics for a single topic in the ["Collection Summary"](#) display.

<b>Connection</b>	The name of the connection.
<b>Host</b>	The name of the host.
<b>Database</b>	The name of the database.
<b>Collection</b>	The name of the collection.
<b>Alert Level</b>	The current alert level. <ul style="list-style-type: none"> <li>● -- One or more alerts have exceeded their specified <b>ALARMLEVEL</b> threshold.</li> <li>● -- One or more alerts have exceeded their specified <b>WARNINGLEVEL</b> threshold.</li> <li>● -- No alerts have exceeded an alert threshold.</li> </ul>
<b>Alert Count</b>	The total number of alerts for the database.
<b>Number of Objects</b>	The total number of objects or documents in the collection.
<b>Average Object Size</b>	The average size, in bytes, of the objects in the collection.
<b>Indexes</b>	The total number of indexes in the collection.

**Expired**

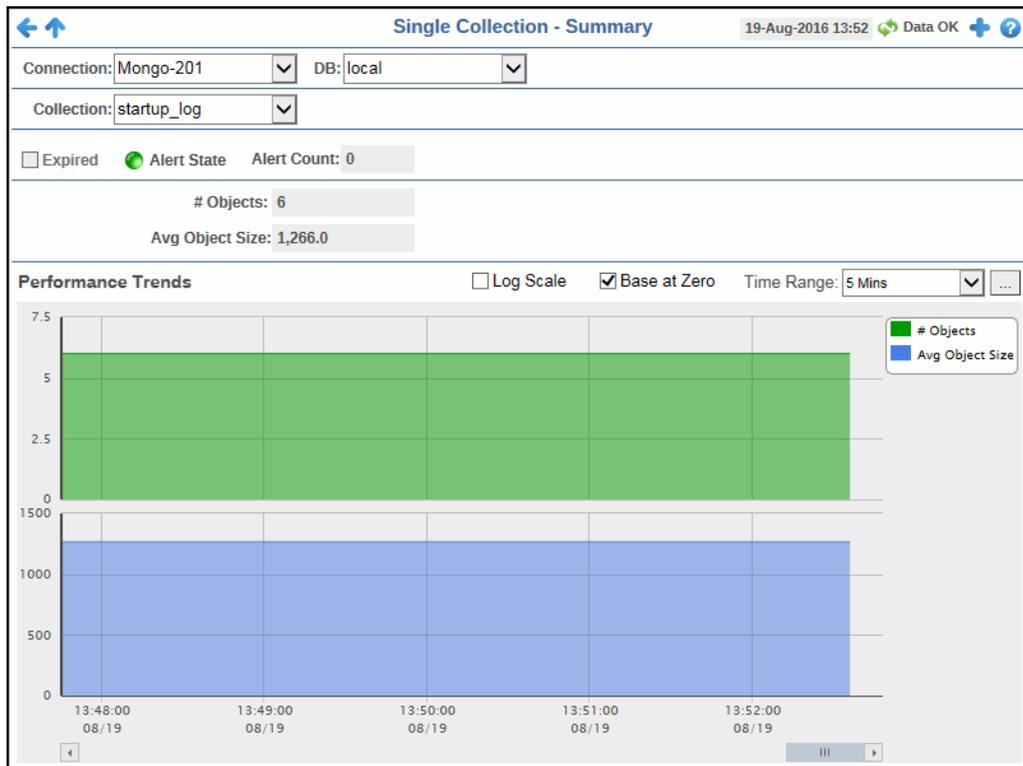
When checked, performance data has not been received within the time specified (in seconds) in the **Expire Time** field in the **Duration** region in the RTView Configuration Application > **(Project Name)** > **Solution Package Configuration** > **MongoDB** > **DATA STORAGE** tab. The **Delete Time** field (also in the **Duration** region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.

**time\_stamp**

The date and time this row of data was last updated.

**Collection Summary**

Track performance and utilization metrics for a single collection on a single database.



**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- 6,047** The number of items currently in the display.

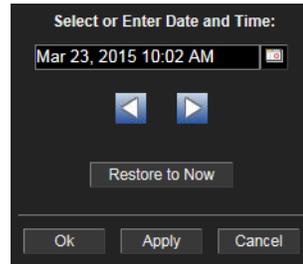
- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

**Fields and Data**

This display includes:

<b>Connection</b>	Select the connection for which you want to view collection data.
<b>DB</b>	Select the database for which you want to view collection data, or select <b>All Databases</b> to view all collections for all databases.
<b>Collection</b>	Select the connection for which you want to view data.
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > ( <b>Project Name</b> ) > <b>Solution Package Configuration</b> > <b>MongoDB</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Alert State</b>	The current alert level.  -- One or more alerts have exceeded their specified <b>ALARMLEVEL</b> threshold.  -- One or more alerts have exceeded their specified <b>WARNINGLEVEL</b> threshold.  -- No alerts have exceeded an alert threshold.
<b>Alert Count</b>	The total number of alerts for the database.
<b># Objects</b>	The total number of objects in the collection.
<b>Avg Object Size</b>	The average size, in bytes, of the objects in the collection.
<b>Performance Trends Graph</b>	Shows message data for the selected collection. <b># Objects</b> -- Traces the total number of objects in the collection. <b>Avg Object Size</b> -- Traces the average size of objects in the collection.
<b>Log Scale</b>	This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.

- Base at Zero** When this option is checked, zero is set as the Y axis minimum for all graph traces.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

---

## MongoDB - HTML

This section describes the HTML version of the Solution Package for MongoDB which features an overview display, "[MongoDB Overview - HTML](#)" (shown below) and the following Views which can be found under **Components** tab > **Databases** > **MongoDB Databases**. For additional details, see vendor documentation.

- ["Mongo Instances View - HTML"](#)
- ["Mongo Databases View - HTML"](#)
- ["Mongo Databases View - HTML"](#)

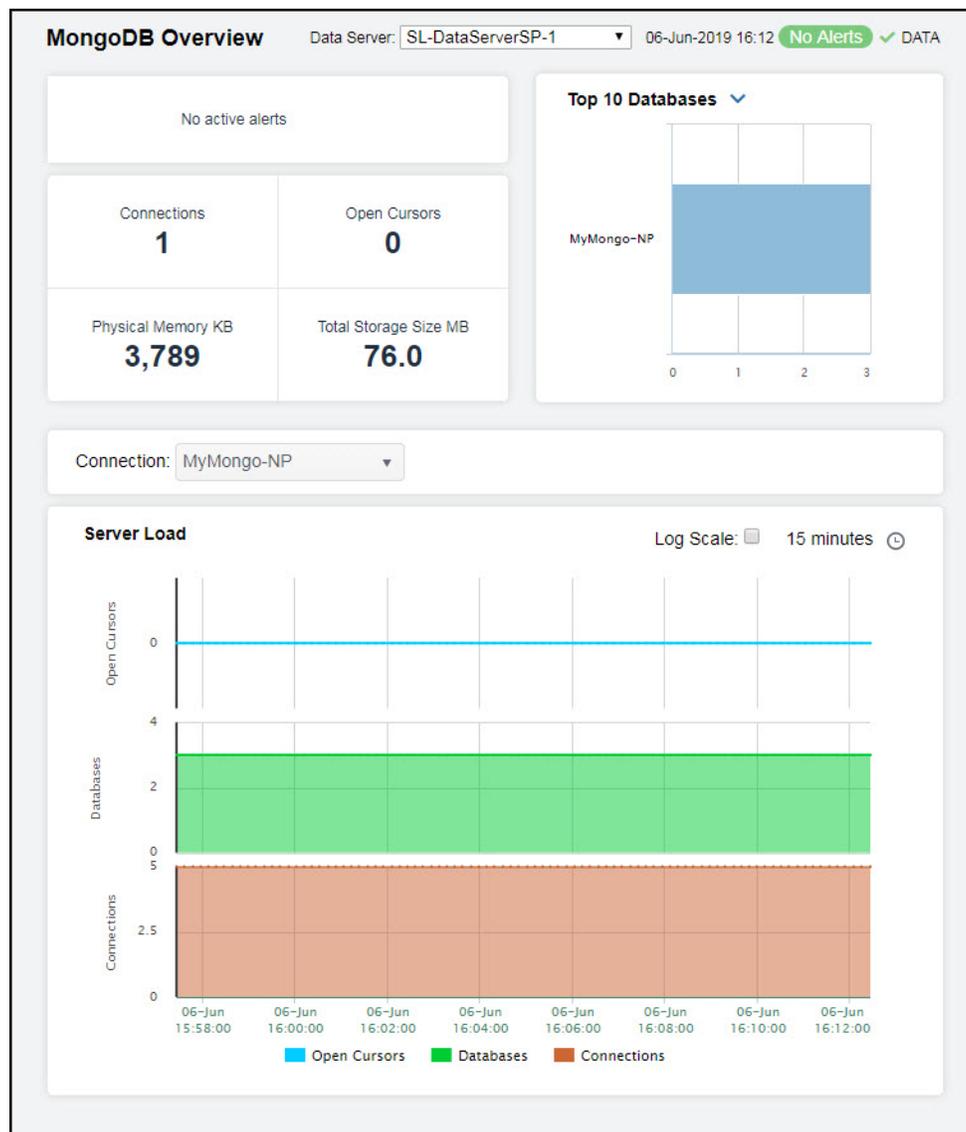
## MongoDB Overview - HTML

The **MongoDB Overview** is the top-level display for the MongoDB Monitor, which provides a good starting point for immediately getting the status of all your databases on your Data Server. You can select the RTView DataServer for which you want to see data and easily view the current data for that DataServer including:

- The total number of active alerts for the selected DataServer, including the total number of critical and warning alerts.
- The total number of connections.
- The total number of open cursors across all instances.
- The total amount of physical memory across all instances.
- The total number of running hosts and the total number of engines.
- The total amount of storage size across all instances.
- A visual list of the top 10 databases based on average back flush, physical memory, open cursors, storage size, free space, and collections on your connected DataServer.

You can hover over each region in the upper half of the Overview to see more detail. You can also drill down to see even more detail by clicking on each respective region in the Overview. For example, clicking on the alerts in the **CRITICAL** and **WARNING** alerts region opens the **Alerts Table by Components** display.

The bottom half of the display provides a trend graph for a selected connection. You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.



## Mongo Instances View - HTML

These displays present performance metrics and alert statuses for all MongoDB instances. Clicking **Mongo Instances** from the left/navigation menu opens the ["MongoDB Instances Table - HTML"](#) display, which shows all available utilization metrics for all MongoDB instances. The following displays are available:

- **All Instances Heatmap:** Opens the ["MongoDB Instances Heatmap - HTML"](#) display, which shows status and alerts for all MongoDB instances in a heatmap format.
- **Single Instance Summary:** Opens the ["MongoDB Instance Summary - HTML"](#) display, which enables you to view available utilization metrics for a single MongoDB instance.

## MongoDB Instances Table - HTML

This display enables you to investigate detailed utilization metrics for all MongoDB Instances. This table contains all metrics available for instances, including the number of current connections. Each row in the table contains data for a particular instance. Click a column header to sort column data in numerical or alphabetical order. Double-click on a table row to drill-down to the ["MongoDB Instance Summary - HTML"](#) display and view metrics for that particular instance. Toggle between the commonly accessed displays by clicking the drop down list on the display title.

Connection	Host	Connected Status	Alert Level	Alert Count	Average Back Flush	C
MyMongo-NP	rhel7vm				0.0	

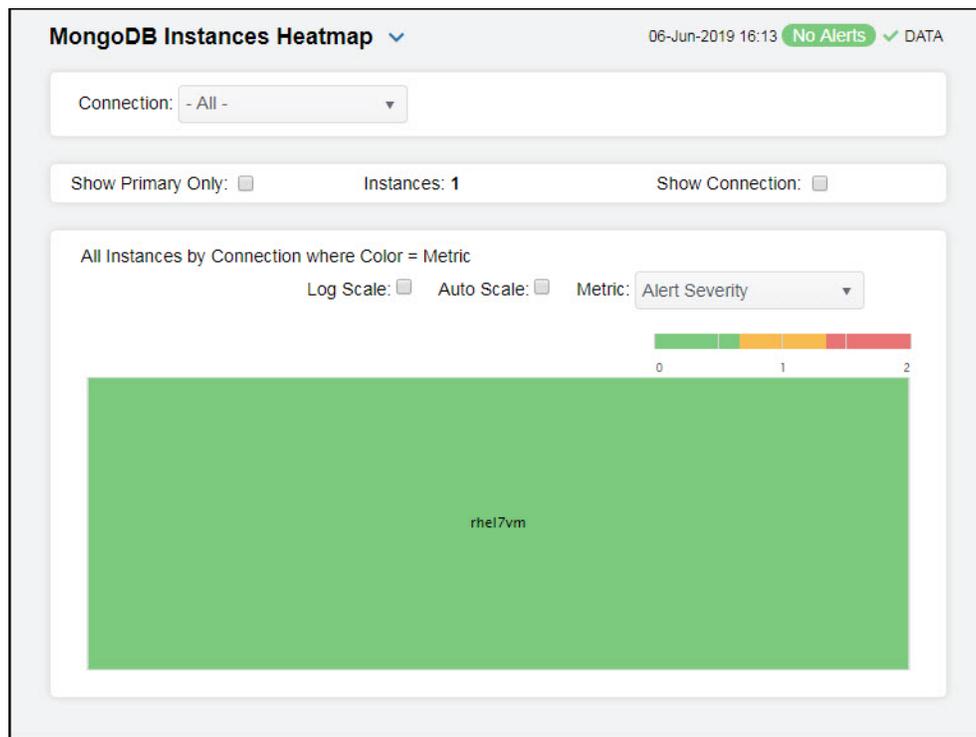
## MongoDB Instances Heatmap - HTML

Clicking **All Instances Heatmap** in the left/navigation menu opens the **MongoDB Instances Heatmap**, which provides an easy-to-view interface that allows you to view the status and alerts of all MongoDB Instances. Use the **Metric** drop-down menu to view the **Alert Severity, Alert Count, Physical Memory, Open Cursors, Connections, or Databases**.

The heatmap is organized by host, each rectangle representing a connection. The rectangle color indicates the most critical alert state. Click on a node to drill-down to the ["MongoDB Instance Summary - HTML"](#) display and view metrics for a particular connection.

You can select the **Primary Only** check box to display connections in the heatmap that have **Designation in Set** (within a replica set) defined as **Primary**, as well as those connections that are not part of a replica set (do not have a defined **Designation in Set**). Those connections with **Designation in Set** defined as **Secondary** will not be displayed. You can select **Show Connection** to display the associated connection in each rectangle.

You can toggle between the commonly accessed displays by clicking the drop down list on the display title. Mouse-over rectangles to view more details about host performance and status.



### Available Metrics

Select the metric driving the heatmap display. The default is **Alert Severity**. Each **Metric** has a color gradient bar that maps values to colors. The heatmap organizes the instances by host, where each rectangle represents an instance. Mouse-over any rectangle to display the current values of the metrics for the instance. Click on a rectangle to drill-down to the associated ["MongoDB Instance Summary - HTML"](#) display for a detailed view of metrics for that particular instance.

#### Alert Severity

The maximum alert level in the item (index) associated with the rectangle. Values range from **0** to **2**, as indicated in the color gradient bar , where **2** is the greatest **Alert Severity**.

**2** -- Metrics that have exceeded their specified **ALARMLEVEL** threshold and have an Alert Severity value of **2** are shown in red. For a given rectangle, this indicates that one or more metrics have exceeded their alarm threshold.

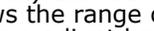
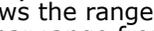
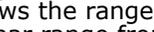
**1** -- Metrics that have exceeded their specified **WARNINGLEVEL** threshold and have an Alert Severity value of **1** are shown in yellow. For a given rectangle, this indicates that one or more metrics have exceeded their warning threshold.

**0** -- Metrics that have not exceeded either specified threshold have an Alert Severity value of **0** and are shown in green. For a given rectangle, this indicates that no metrics have exceeded a specified alert threshold.

#### Alert Count

The total number of alarm and warning alerts in a given item (index) associated with the rectangle.

The color gradient bar shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

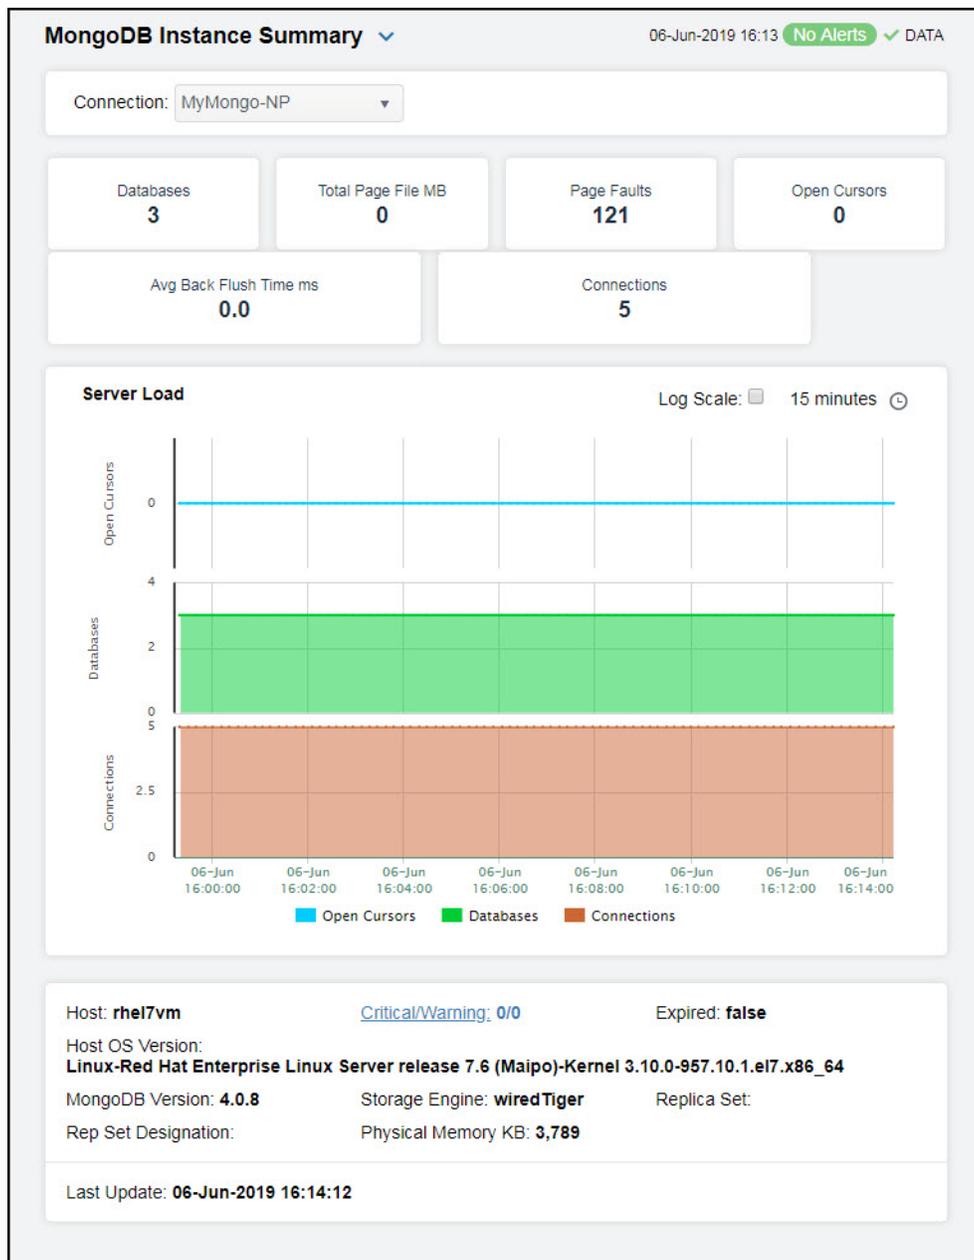
- Physical Memory** The total amount of physical memory currently being used in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum amount of physical memory in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
- The **Auto** option does not impact this metric.
- Open Cursors** The total number of open cursors in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of **MongoInstanceOpenCursorsHigh**, which is **2000**. The middle value in the gradient bar indicates the middle value of the range (the default is **1000**).
- When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.
- Connections** The total number of connections in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of connections in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
- The **Auto** option does not impact this metric.
- Databases** The total number of databases in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of databases in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
- The **Auto** option does not impact this metric.

## MongoDB Instance Summary - HTML

Clicking **Single Instance Summary** in the left/navigation menu opens the **MongoDB Instance Summary** display, which allows you to view current as well as trending data for specific instances. Clicking on the information boxes at the top of the display takes you to the ["MongoDB Instances Table - HTML"](#) display, where you can view additional instance data.

The **Server Load** trend graph allows you to view trend data for the open cursors, databases, and connections over a selected time range.

Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display.



## Mongo Databases View - HTML

These displays present detailed performance metrics and alert statuses for all databases (in a heatmap or a tabular format) or for an individual database. Clicking **Mongo Databases** from the left/navigation menu opens the ["MongoDB Databases Table - HTML"](#) display, which shows a tabular view of all databases their associated metrics. The following displays are available:

- **All Databases Heatmap:** Opens the ["MongoDB Databases Heatmap - HTML"](#) display, which shows status and alerts for all MongoDB databases in a heatmap format.
- **Single Database Summary:** Opens the ["MongoDB Single Database Summary - HTML"](#) display, which enables you to view available utilization metrics for a single MongoDB database.

**Note:** No database information will display in the heatmap, table, or summary displays if a connection cannot be established.

### MongoDB Databases Table - HTML

View details for all databases in a single server. Each row in the table contains data for a particular database. Click a column header to sort column data in numerical or alphabetical order. Double-click on a table row to drill-down to the ["MongoDB Single Database Summary - HTML"](#) display and view metrics for that particular database. Toggle between the commonly accessed displays by clicking the drop down list on the display title.

**MONGO Databases Table** 06-Jun-2019 16:14 No Alerts ✓ DATA

Connection: MyMongo-NP

Databases: **3**

**All Databases Table**

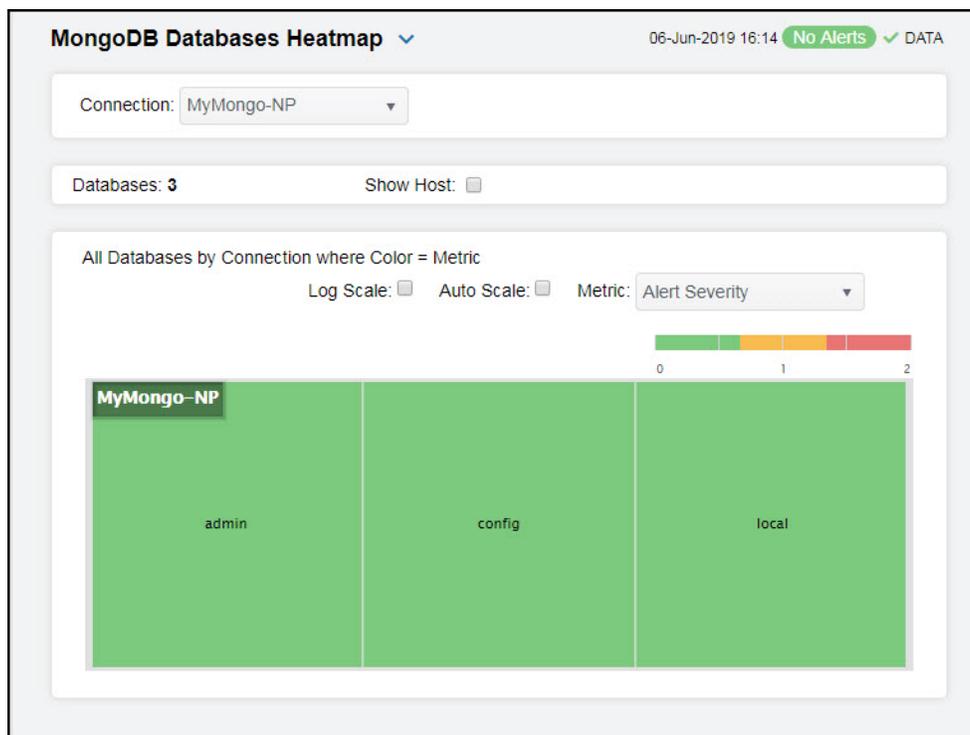
Connection	Host	Name	Alert Level	Alert Count	Status Up Down
MyMongo-NP	rhel7vm	admin	✓		✓
MyMongo-NP	rhel7vm	config	✓		✓
MyMongo-NP	rhel7vm	local	✓		✓

## MongoDB Databases Heatmap - HTML

Clicking **All Databases Heatmap** in the left/navigation menu opens the **MongoDB Databases Heatmap**, which provides an easy-to-view interface that allows you to track utilization and performance metrics for all databases in a heatmap format. Use the **Metric** drop-down menu to view the heatmap using the alert severity, alert count, collections, or data size metrics.

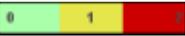
The heatmap is organized by host, each rectangle representing a database. The rectangle color indicates the most critical alert state. Click on a node to drill-down to the ["MongoDB Single Database Summary - HTML"](#) display and view metrics for a particular database.

You can select the **Show Host** check box to display name of the host in each rectangle. You can toggle between the commonly accessed displays by clicking the drop down list on the display title. Mouse-over rectangles to view more details about host performance and status.



### Available Metrics

Select the metric driving the heatmap display. The default is **Alert Severity**. Each **Metric** has a color gradient bar that maps values to colors. The heatmap organizes the databases by connection, where each rectangle represents a database. Mouse-over any rectangle to display the current values of the metrics for the database. Click on a rectangle to drill-down to the associated ["MongoDB Single Database Summary - HTML"](#) display for a detailed view of metrics for that particular database.

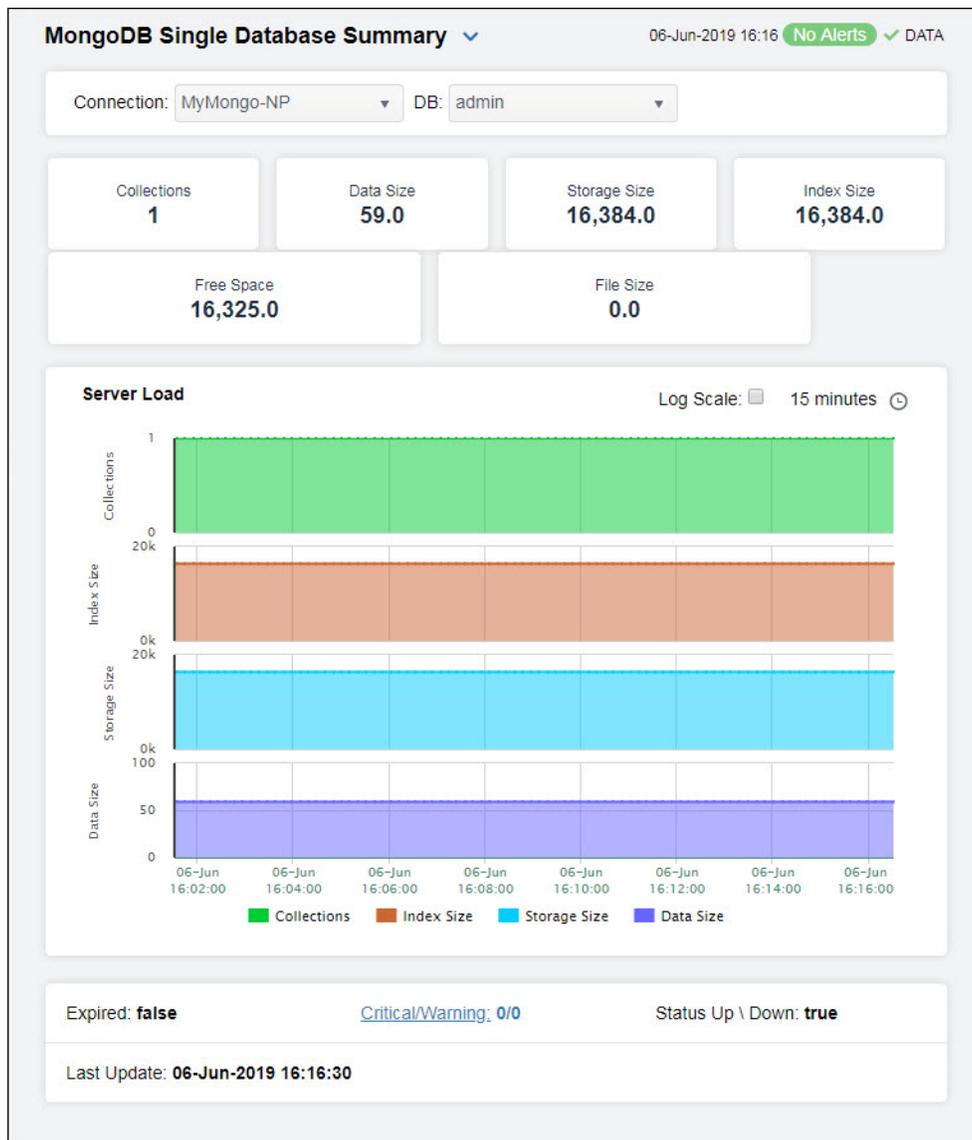
- Alert Severity** The maximum alert level in the item (index) associated with the rectangle. Values range from **0** to **2**, as indicated in the color gradient bar , where **2** is the greatest **Alert Severity**.
- 2** -- Metrics that have exceeded their specified **ALARMLEVEL** threshold and have an Alert Severity value of **2** are shown in red. For a given rectangle, this indicates that one or more metrics have exceeded their alarm threshold.
- 1** -- Metrics that have exceeded their specified **WARNINGLEVEL** threshold and have an Alert Severity value of **1** are shown in yellow. For a given rectangle, this indicates that one or more metrics have exceeded their warning threshold.
- 0** -- Metrics that have not exceeded either specified threshold have an Alert Severity value of **0** and are shown in green. For a given rectangle, this indicates that no metrics have exceeded a specified alert threshold.
- Alert Count** The total number of alarm and warning alerts in a given item (index) associated with the rectangle.
- The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
- Collections** The total number of collections in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum number of collections in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
- The **Auto Scale** option does not impact this metric.
- Data Size** The total size (in bytes) of the data in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of **MongoDatabaseDataSizeHigh**, which is **100,000**. The middle value in the gradient bar indicates the middle value of the range (the default is **50,000**).
- When **Auto Scale** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

## MongoDB Single Database Summary - HTML

Clicking **Single Database Summary** in the left/navigation menu opens the **MongoDB Database Summary** display, which allows you to view utilization, performance, and trend data for a specific database. Clicking on the information boxes at the top of the display takes you to the ["MongoDB Databases Table - HTML"](#) display, where you can view additional database data.

The **Performance Trends** trend graph allows you to view trend data for the collections, index size, storage size, and data size over a selected time range.

Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display.



## Mongo Collections View - HTML

These displays present several views of performance metrics for collections. Clicking **Mongo Collections** from the left/navigation menu opens the ["MongoDB Collections Table - HTML"](#) display, which allows you to view performance and utilization metrics for all collections in a particular database, or for all collections on all databases. The following displays are available:

- **All Collections Heatmap:** Opens the ["MongoDB Collections Heatmap - HTML"](#) display, which allows you to view performance and utilization metrics for all collections that exist in each of your connections.
- **Single Collection Summary:** Opens the ["MongoDB Single Collection Summary - HTML"](#) display, which shows detailed performance and utilization metrics and trends for a specified collection on a particular database.

### MongoDB Collections Table - HTML

Track performance and utilization metrics for all collections on a single database, or for all connections on all databases. Click a column header to sort column data in numerical or alphabetical order. Double-click on a table row to drill-down to the ["MongoDB Single Collection Summary - HTML"](#) display and view metrics for that particular collection. Toggle between the commonly accessed displays by clicking the drop down list on the display title.

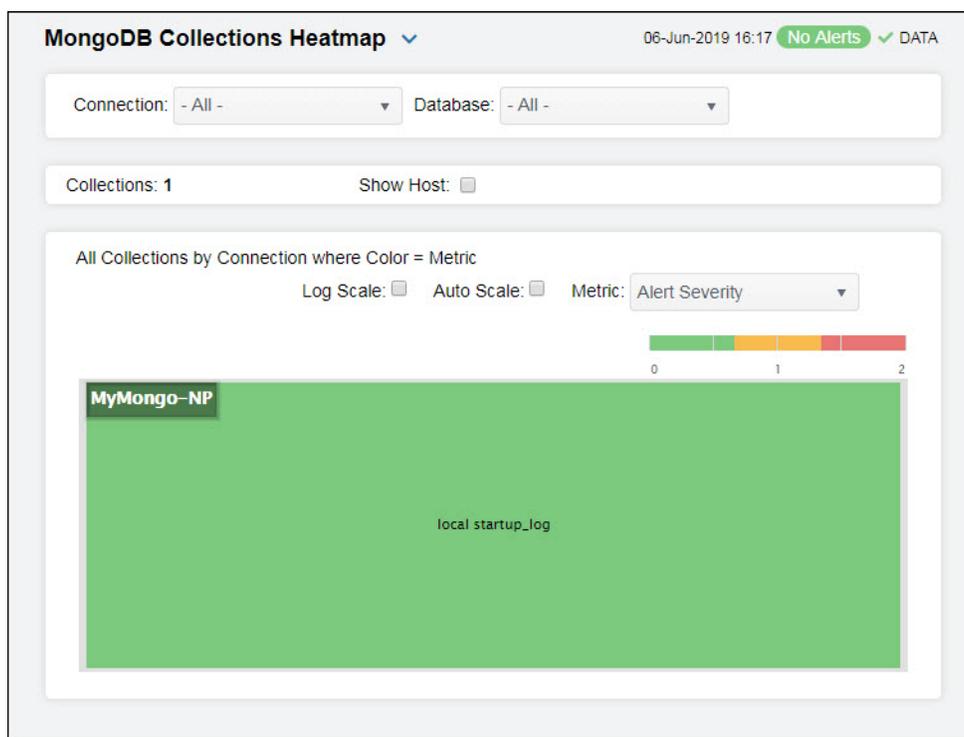
Connection	Host	Database	Name	Alert Level	Alert Count
MyMongo-NP	rhel7vm	local	startup_log	✓	

## MongoDB Collections Heatmap - HTML

Clicking **All Collections Heatmap** in the left/navigation menu opens the **MongoDB Collections Heatmap**, which provides a heatmap view of the status and alerts of all collections within each connection. Use the **Metric** drop-down menu to view the heatmap using alert severity, alert count, number of objects, or average object size metrics.

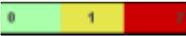
The heatmap is organized by connection, with each rectangle representing a collection. The rectangle color indicates the most critical alert state. Click on a node to drill-down to the ["MongoDB Single Collection Summary - HTML"](#) display and view metrics for a particular collection.

You can select the **Show Host** check box to display name of the host in each rectangle. You can toggle between the commonly accessed displays by clicking the drop down list on the display title. Mouse-over rectangles to view more details about host performance and status.



### Available Metrics

Select the metric driving the heatmap display. The default is **Alert Severity**. Each **Metric** has a color gradient bar that maps values to colors. The heatmap organizes the collections by connection, where each rectangle represents a collection. Mouse-over any rectangle to display the current values of the metrics for the collection. Click on a rectangle to drill-down to the associated ["MongoDB Single Collection Summary - HTML"](#) display for a detailed view of metrics for that particular collection.

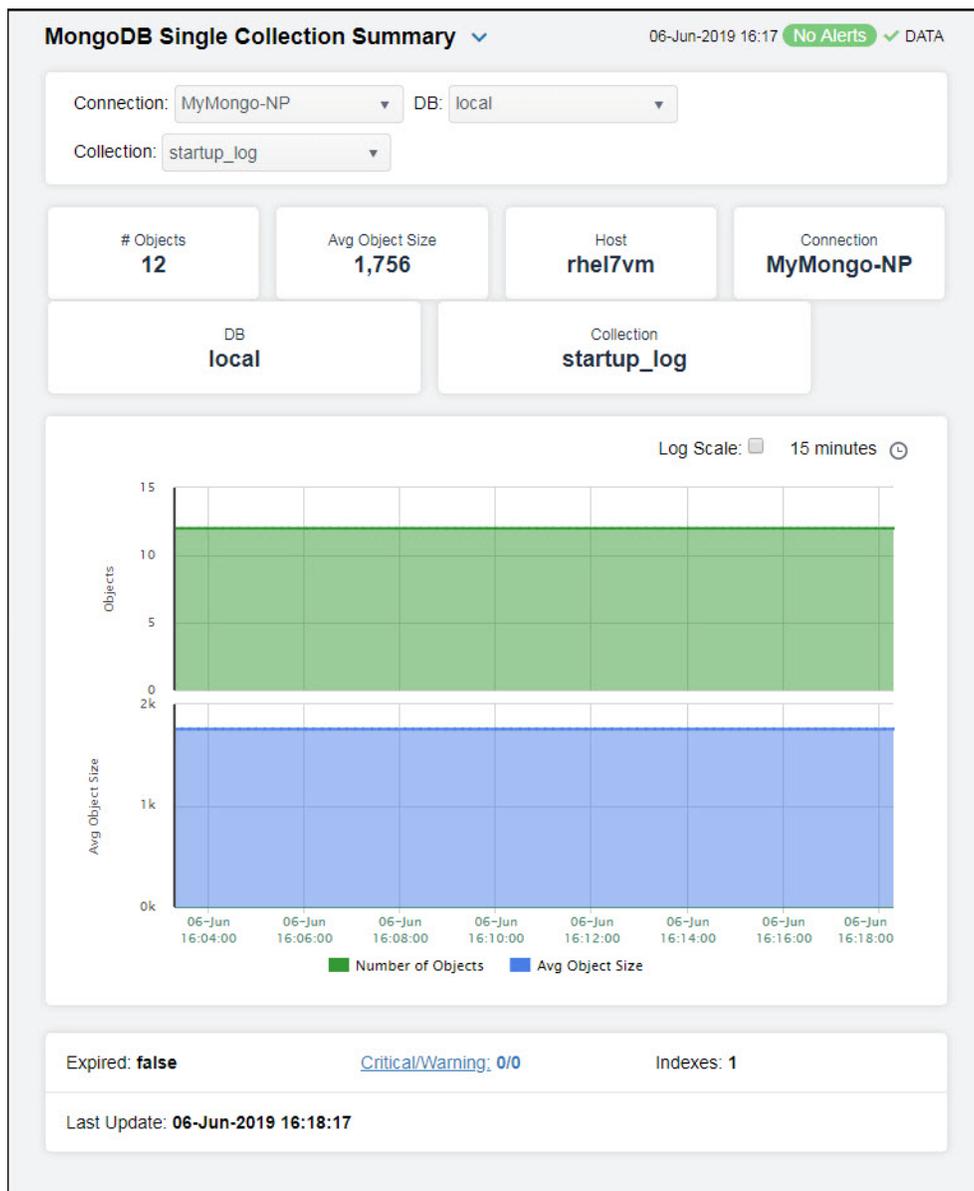
<b>Alert Severity</b>	<p>The maximum alert level in the item (index) associated with the rectangle. Values range from <b>0</b> to <b>2</b>, as indicated in the color gradient bar , where <b>2</b> is the greatest <b>Alert Severity</b>.</p> <p><b>2</b> -- Metrics that have exceeded their specified <b>ALARMLEVEL</b> threshold and have an Alert Severity value of <b>2</b> are shown in red. For a given rectangle, this indicates that one or more metrics have exceeded their alarm threshold.</p> <p><b>1</b> -- Metrics that have exceeded their specified <b>WARNINGLEVEL</b> threshold and have an Alert Severity value of <b>1</b> are shown in yellow. For a given rectangle, this indicates that one or more metrics have exceeded their warning threshold.</p> <p><b>0</b> -- Metrics that have not exceeded either specified threshold have an Alert Severity value of <b>0</b> and are shown in green. For a given rectangle, this indicates that no metrics have exceeded a specified alert threshold.</p>
<b>Alert Count</b>	<p>The total number of alarm and warning alerts in a given item (index) associated with the rectangle.</p> <p>The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>Objects</b>	<p>The total number of objects or documents in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from <b>0</b> to the alert threshold of <b>MongoCollectionNumObjectsHigh</b>. The middle value in the gradient bar indicates the middle value of the range.</p> <p>When <b>Auto Scale</b> is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>
<b>Avg Object Size</b>	<p>The average size (in bytes) of an object in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of objects in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>

## MongoDB Single Collection Summary - HTML

Clicking **Single Collection Summary** in the left/navigation menu opens the **MongoDB Collection Summary** display, which allows you to track performance and utilization metrics for a single collection on a single database. Clicking on the information boxes at the top of the display takes you to the ["MongoDB Collections Table - HTML"](#) display, where you can view additional database data.

The trend graph allows you to view trend data for the number of objects and average object size over a selected time range.

Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display.



---

## MySQL Database

The MySQL Databases displays provide extensive visibility into the health and performance of the MySQL database. For information about the HTML version, see ["MySQL Database - HTML"](#).

The following MySQL Database Views (and their associated displays) can be found under **Components** tab > **Databases** > **MySQL Database**:

- ["All MySQL Databases"](#)
- ["Single MySQL Database"](#)

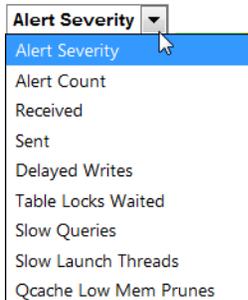
### All MySQL Databases

Displays in this View are:

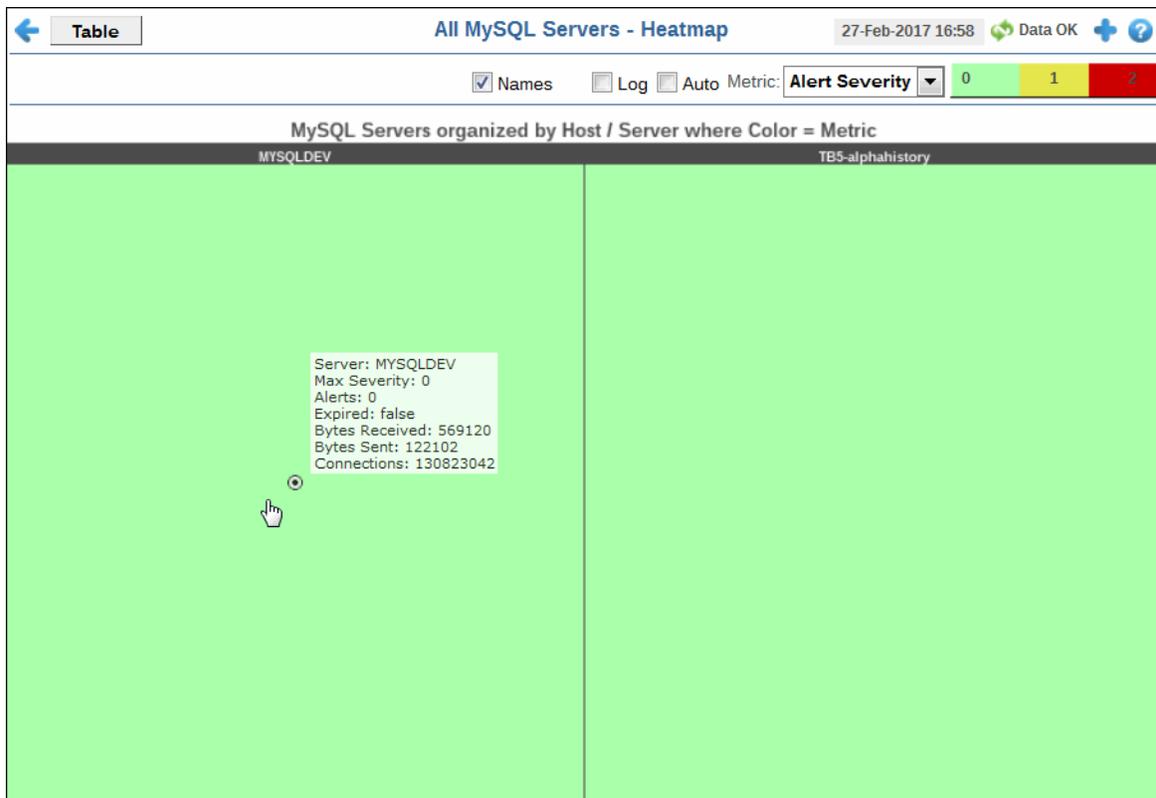
- ["All Servers Heatmap"](#): A heatmap view of all servers and their associated metrics.
- ["All Servers Table"](#): A tabular view of your servers and their associated metrics.

### All Servers Heatmap

This heatmap display provides an easy-to-view interface that allows you to quickly identify the current status of each of your servers. Choose a metric from the **Metric** drop down menu. By default, this display shows the heatmap based on the **Alert Severity** metric. Other metrics are Alert Count, Received, Sent, Delayed Writes, Table Locks Waited, Slow Queries, Slow Launch Threads and Qcache Low Mem Prunes.



Each rectangle in the heatmap is a different server. Use the **Names** check-box  to include or exclude labels in the heatmap, and mouse over a rectangle to see additional metrics for a server. Click a rectangle to open the “[Server Summary](#)” display and see additional details for the selected server.



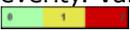
#### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

- Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

#### Fields and Data:

- Names** Select this check box to display the names of the instances at the top of each rectangle in the heatmap.
- Log** Select to this check box to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Auto** Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value.  
**Note:** Some metrics auto-scale automatically, even when **Auto** is not selected.

<b>Metric</b>	Choose a metric to view in the display. For details about the data, refer to vendor documentation.
<b>Alert Severity</b>	<p>The current alert severity. Values range from <b>0</b> - <b>2</b>, as indicated in the color gradient  bar, where <b>2</b> is the highest Alert Severity:</p> <ul style="list-style-type: none"> <li><span style="color: red;">●</span> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li><span style="color: yellow;">●</span> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li><span style="color: green;">●</span> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Alert Count</b>	<p>The total number of critical and warning unacknowledged alerts. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.</p>
<b>Received</b>	<p>The total number of bytes received. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the alarm threshold specified for the <b>MysqlBytesReceivedHigh</b> alert. The middle value in the gradient bar indicates the average count.</p>
<b>Sent</b>	<p>The total number of bytes sent. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the alarm threshold specified for the <b>MysqlBytesSentHigh</b> alert. The middle value in the gradient bar indicates the average count.</p>
<b>Delayed Writes</b>	<p>The total number of delayed writes. Values range from <b>0</b> to the alarm threshold specified for the <b>MysqlDelayedWrites</b> alert. The middle value in the gradient bar indicates the average count:</p> <ul style="list-style-type: none"> <li><span style="color: red;">●</span> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li><span style="color: yellow;">●</span> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li><span style="color: green;">●</span> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Table Locks Waited</b>	<p>The total number of table locks waited. Values range from <b>0</b> to the alarm threshold specified for the <b>MysqlLocksWaited</b> alert. The middle value in the gradient bar indicates the average count:</p> <ul style="list-style-type: none"> <li><span style="color: red;">●</span> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li><span style="color: yellow;">●</span> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li><span style="color: green;">●</span> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Slow Queries</b>	<p>The total number of slow queries. Values range from <b>0</b> to the alarm threshold specified for the <b>MysqlSlowQueries</b>. The middle value in the gradient bar indicates the average count:</p> <ul style="list-style-type: none"> <li><span style="color: red;">●</span> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li><span style="color: yellow;">●</span> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li><span style="color: green;">●</span> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>

**Slow Launch Threads**

The total number of slow launch threads. Values range from **0** to the alarm threshold specified for the **MySqlSlowThreads**. The middle value in the gradient bar indicates the average count:

-  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
-  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
-  Green indicates that no metrics have exceeded their alert thresholds.

**Qcache Low Mem Prunes**

The total number of Qcache low memory prunes. Values range from **0** to the alarm threshold specified for the **MySqlQcacheLowMemPrunes**. The middle value in the gradient bar indicates the average count:

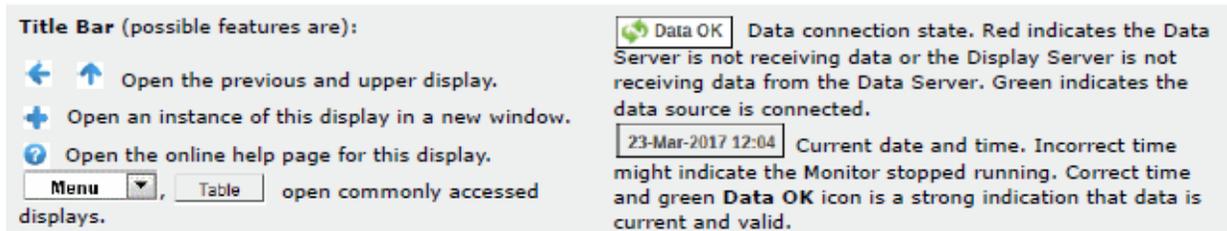
-  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
-  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
-  Green indicates that no metrics have exceeded their alert thresholds.

## All Servers Table

This display provides a tabular view of the performance metrics shown in the “[All Servers Heatmap](#)” (alert level, alert count, bytes received, and so forth), as well as additional metrics (such as query information and uptime).

Each table row is a different server. Click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by clicking a row to view details for a server in the “[Server Summary](#)” display.

Heatmap										
All MySQL Servers Table										
27-Feb-2017 17:04 Data OK										
Server Name	Expired	Alert Level	Alert Count	Connected	Last Query	Avg Exec Time	Avg Process Time	Bytes Received	Bytes Sent	
MYSQLDEV	<input type="checkbox"/>		0	<input checked="" type="checkbox"/>	OK	0.24	0.24	425,250	468	



**All MySQL Servers Table**

<b>Server Name</b>	The name of the server.
<b>Expired</b>	<p>When checked, performance data about the server has not been received within the time specified (in seconds) in the <b>\$mysqlRowExpirationTime</b> field in the <b>conf\rtvapm_mysqlmon.properties</b> file. The <b>\$mysqlRowExpirationTimeForDelete</b> field allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response from the server. To view/edit the current values, modify the following lines in the <b>.properties</b> file:</p> <pre>##### # CACHE / HISTORIAN SETTINGS # collector.sl.rtvview.sub=\$mssqlRowExpirationTime:120 collector.sl.rtvview.sub=\$mssqlRowExpirationTimeForDelete:0</pre> <p>In the example above, the <b>Expired</b> check box would be checked after 120 seconds, and the row would never be deleted. If <b>\$mysqlRowExpirationTimeForDelete</b> was set to 3600, then the row would be removed from the table after 3600 seconds.</p>
<b>Alert Level</b>	<p>The current alert severity.</p> <ul style="list-style-type: none"> <li><span style="color: red;">●</span> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li><span style="color: yellow;">●</span> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li><span style="color: green;">●</span> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Alert Count</b>	The total number of alerts for the server.
<b>Connected</b>	When checked, the server is connected.
<b>Last Query</b>	The status of the last query made:
<b>Avg Exec Time</b>	The average amount of execution time, in seconds.
<b>Avg Process Time</b>	The average amount of process time, in seconds.
<b>Bytes Received</b>	The total number of bytes received since the server was last started.
<b>Connections</b>	The total number of connections since the server was last started.
<b>Delayed Writes</b>	The total number of delayed writes.
<b>Queries</b>	The total number of queries.
<b>Query Objects</b>	The total number of query objects.
<b>Slow Queries</b>	The total number of slow queries.

<b>Total Executions</b>	The total number of executions.
<b>Uptime</b>	The amount of time since the server was last started, in seconds.
<b>Concurrent</b>	When checked, the database allows concurrent usage.
<b>Enabled</b>	When checked, the database is enabled for usage.
<b>Timestamp</b>	The data and time of the last data update.

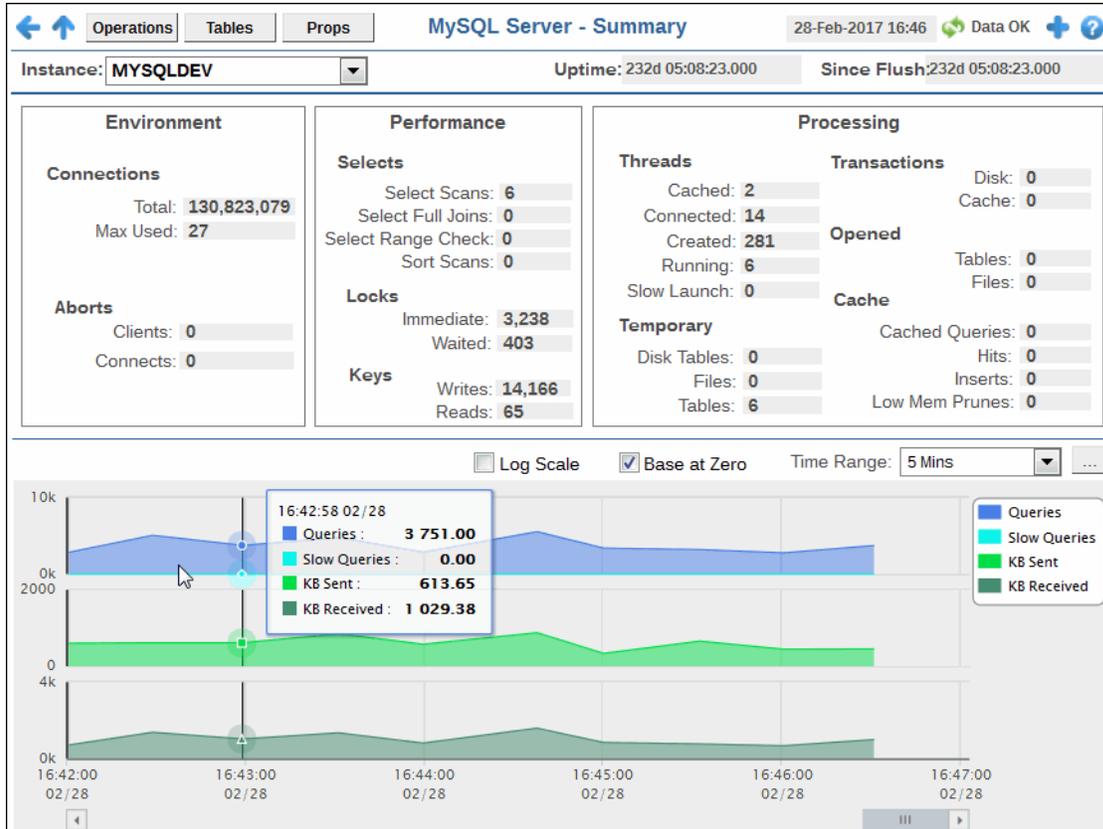
## Single MySQL Database

Displays in this View are:

- **"Server Summary"**: Displays performance, processing, alerts, memory, and trend data for a particular database server.
- **"Servers Properties"**: Displays the values of properties on servers.
- **"Servers Operations"**: Trend graph that traces server queries, slow queries, KB sent and KB received.
- **"User Tables"**: A tabular view of cache tables performance and utilization metrics.

## Server Summary

View connection, performance and processing details for a single MySQL database server, as well as trending data for the number of kilobytes received and queries. Choose an instance from the **Instance** drop-down menu. Mouse over the trend graph to see performance metrics with time stamps.



**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Filter By:**

**Instance:** Select the instance for which you want to show data in the display.

**Fields and Data:** For details about the data in this display, please refer to vendor documentation.

**Uptime** The amount of time since the server was last started, in number of days, hours, minutes and seconds.

**Since Flush** The amount of time since the last flush, in number of days, hours, minutes and seconds.

**Performance Trend Graph**

Traces the following:

**Queries:** Traces the amount queries per second.

**Slow Queries:** Traces the amount of slow queries per second.

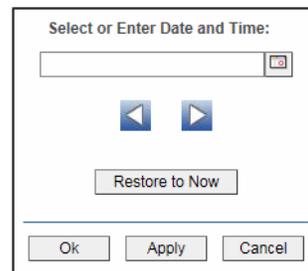
**KB Sent:** Traces the number of kilobytes sent per second.

**KB Received:** Traces the number of kilobytes received per second.

**Log** Select to this check box to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.

**Base at Zero** Select to use zero (0) as the Y axis minimum for all graph traces.

**Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Servers Properties

View properties and property values for a single MySQL database server.

Choose an instance from the **Instance** drop-down menu. Each table row is a different property for the selected instance. Enter a search string in the **Property Filter** field to limit the number of table rows. Click a column header to sort column data in numerical or alphabetical order.

The screenshot shows the 'MySQL Server - Properties' window. At the top, there are navigation buttons (back, forward, home) and tabs for 'Summary', 'Operations', and 'Tables'. The title bar includes the instance name 'MySQLDEV', the uptime '231d 05:33:04.000', and the 'Since Flush' time '231d 05:33:04.000'. A 'Property Filter' field is present, and the 'Last Update' is '27-Feb-2017 05:06'. The main area contains a table with two columns: 'Property' and 'Value'. The table lists various MySQL configuration parameters and their current values. At the bottom, there is a pagination control showing 'Page 1 of 2' and '1 - 200 of 259 items'.

Property	Value
auto_increment_increment	1
auto_increment_offset	1
autocommit	ON
automatic_sp_privileges	ON
back_log	50
basedir	C:\Program Files\MySQL\MySQL Server 5.5\
big_tables	OFF
binlog_cache_size	32768
binlog_direct_non_transactional_updates	OFF
binlog_format	STATEMENT
binlog_stmt_cache_size	32768
bulk_insert_buffer_size	8388608
character_set_client	latin1
character_set_connection	latin1
character_set_database	latin1
character_set_filesystem	binary
character_set_results	latin1
character_set_server	latin1
character_set_system	utf8
character_sets_dir	C:\Program Files\MySQL\MySQL Server 5.5\share\charsets\

**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

### Filter By:

**Instance** Select the database for which you want to show data in the display.

### Fields and Data:

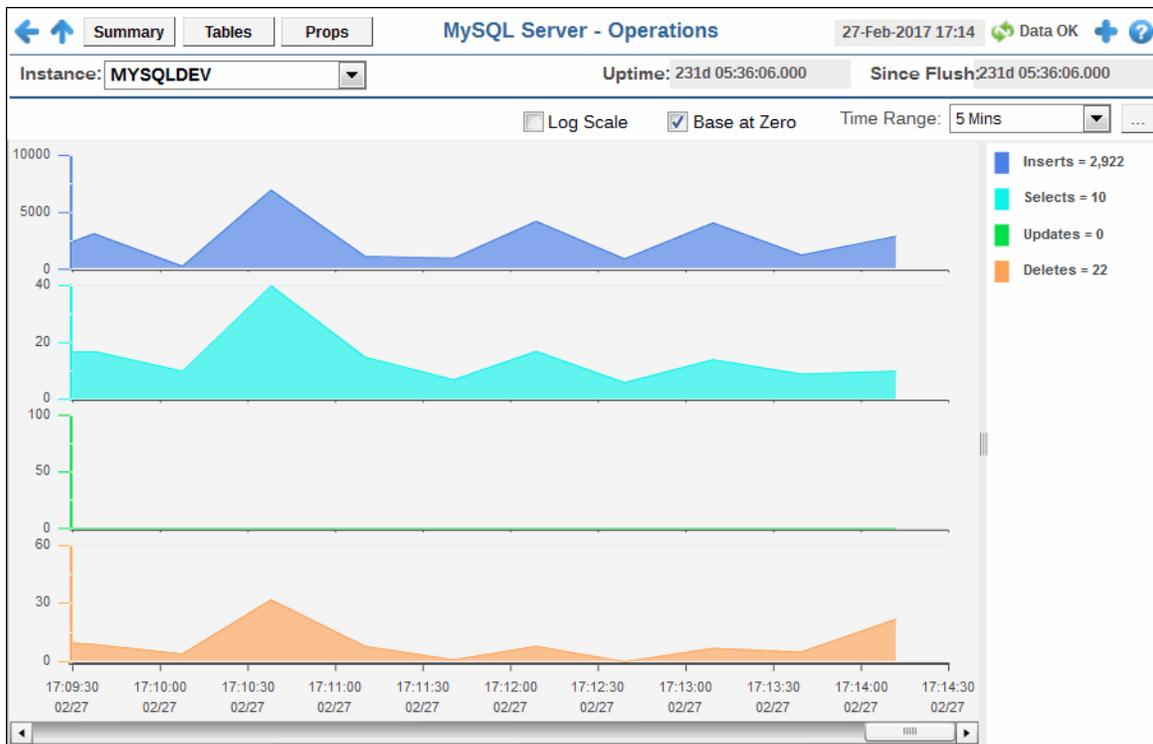
**Uptime** The amount of time since the server was last started, in number of days, hours, minutes and seconds.

**Property Filter:** Enter a search string to filter the number of table rows.

**Since Flush** The amount of time since the last flush, in number of days, hours, minutes and seconds.

## Servers Operations

View trending performance data for a single MySQL database server: **Inserts**, **Selects**, **Updates** and **Deletes**. Choose an instance from the **Instance** drop-down menu. Mouse over the trend graph to see performance metrics with time stamps.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

### Filter By:

**Instance** Select the database for which you want to show data in the display.

### Fields and Data:

**Uptime** The amount of time since the server was last started, in number of days, hours, minutes and seconds.

**Property Filter:** Enter a search string to filter the number of table rows.

**Since Flush** The amount of time since the last flush, in number of days, hours, minutes and seconds.

**Performance Trend Graph**

Traces the following:

**Inserts:** Traces the number of inserts per second.

**Selects:** Traces the number of selects per second.

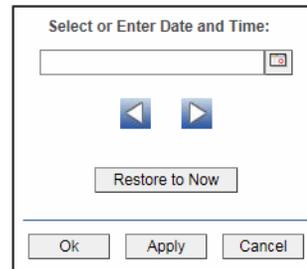
**Updates:** Traces the number of updates per second.

**Deletes:** Traces the number of deletes per second.

**Log** Select to this check box to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.

**Base at Zero** Select to use zero (**0**) as the Y axis minimum for all graph traces.

**Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## User Tables

View performance and utilization details for cache tables for a single MySQL database server. Each row is a different cache table. Choose an instance from the **Instance** drop-down menu. Click a column header to sort column data in numerical or alphabetical order.

Schema	Table	Row Count	Index Size	Data Size	Total Size	Data Free	Engine
alertdefs	alertlevels	0	1,024	0	1,024	0	MyISAM
alertdefs	audit_table	0	1,024	0	1,024	0	MyISAM
rtvhistory	\$bw6_activities_table	515,918	13,483,008	47,221,756	60,704,764	0	MyISAM
rtvhistory	\$bw6_activity_totals_table	56,463	1,383,424	6,107,932	7,491,356	0	MyISAM
rtvhistory	\$bw6_process_totals_app t	9,959	368,640	1,229,296	1,597,936	312,956	MyISAM
rtvhistory	\$bw6_process_totals_appnc	59,718	2,533,376	6,862,184	9,395,560	1,396,252	MyISAM
rtvhistory	\$bw6_process_totals_appsi	9,462	262,144	752,816	1,014,960	0	MyISAM
rtvhistory	\$bw6_process_totals_table	109,461	4,017,152	14,284,080	18,301,232	4,099,164	MyISAM
rtvhistory	\$bw6_processes_table	104,214	2,779,136	8,586,004	11,365,140	0	MyISAM
rtvhistory	bw6_activity_totals	226,128	4,355,072	20,718,016	25,073,088	0	MyISAM
rtvhistory	bw6_appnodes	39,409	764,928	2,597,056	3,361,984	0	MyISAM
rtvhistory	bw6_process_totals	94,395	1,859,584	7,650,588	9,510,172	0	MyISAM
rtvhistory	bw6_process_totals_app	10,979	216,064	777,800	993,864	0	MyISAM
rtvhistory	bw6_process_totals_appnoc	65,919	1,270,784	4,415,924	5,686,708	0	MyISAM
rtvhistory	bw6_process_totals_appsic	65,961	1,274,880	5,211,584	6,486,464	0	MyISAM
rtvhistory	bw6_processes	0	2,048	0	2,048	0	MyISAM
rtvhistory	bw_activities	3,520,325	35,879,936	330,112,152	365,992,088	0	MyISAM
rtvhistory	bw_activity_totals	1,202,835	38,381,568	158,427,548	196,809,116	692,108	MyISAM
rtvhistory	bw_engines	106,159	4,043,776	14,760,112	18,803,888	820,200	MyISAM
rtvhistory	bw_process_totals	78,638	4,087,808	15,453,984	19,541,792	5,266,124	MyISAM
rtvhistory	bw_processes	974,430	39,562,240	198,494,576	238,056,816	47,194,296	MyISAM
rtvhistory	bw_servers	30,982	1,239,040	2,314,796	3,553,836	231,836	MyISAM
rtvhistory	ems_admstats	8,309	158,720	187,194	345,914	12,705	MyISAM
rtvhistory	ems_compedesstotals	270,012	2,754,560	8,640,384	11,394,944	0	MyISAM
rtvhistory	ems_connections	534,561	5,451,776	39,159,128	44,610,904	0	MyISAM
rtvhistory	ems_consumers	2,018,788	20,578,304	87,000,188	117,677,492	0	MyISAM

**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** **Table** open commonly accessed displays.
- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

### Filter By:

**Instance** Select the database for which you want to show data in the display.

**Fields and Data:** For details about the data in this display, please refer to vendor documentation.

**Uptime** The amount of time since the server was last started, in number of days, hours, minutes and seconds.

**Property Filter:** Enter a search string to filter the number of table rows.

**Since Flush** The amount of time since the last flush, in number of days, hours, minutes and seconds.

### Table

<b>Schema</b>	The name of the database.
<b>Table</b>	The name of the table.
<b>Row Count</b>	The number of rows currently in the table.
<b>Index Size</b>	The size of the table indexes, in bytes.
<b>Data Size</b>	The size of the data stored in the table, in bytes (Total Size - Index Size = Data Size).
<b>Total Size</b>	The total size of the table, in bytes.
<b>Data Free RX</b>	The amount of available space that can be reclaimed to store new data, in bytes.
<b>Engine</b>	The storage engine handling the SQL operations.
<b>Last Updated</b>	The time of the last data update.

---

## MySQL Database - HTML

The MySQL Databases HTML displays provide extensive visibility into the health and performance of the MySQL database. The HTML version features an overview display, “[MySQL Overview Display- HTML](#)” (pictured below), and the following Views which can be found under **Components** tab > **Databases** > **MySQL Database**:

- [“MySQL Overview Display- HTML”](#)
- [“All MySQL Instances View - HTML”](#)
- [“Single MySQL Instance View - HTML”](#)

### MySQL Overview Display- HTML

The **MySQL Overview** is the top-level display for the MySQL Solution Package, which provides a good starting point for immediately getting the status of all your MySQL instances on your Data Server.

You can select the RTView DataServer for which you want to see data and easily view the current data for that DataServer including:

- The total number of active alerts, including the total number of critical and warning alerts.
- The number of full joins and scan selects across all servers.
- The number of threads running and threads created across all servers.
- The number of slow queries and total queries across all servers.
- The number of connections.
- A visual list of the top 10 servers with the most queries.

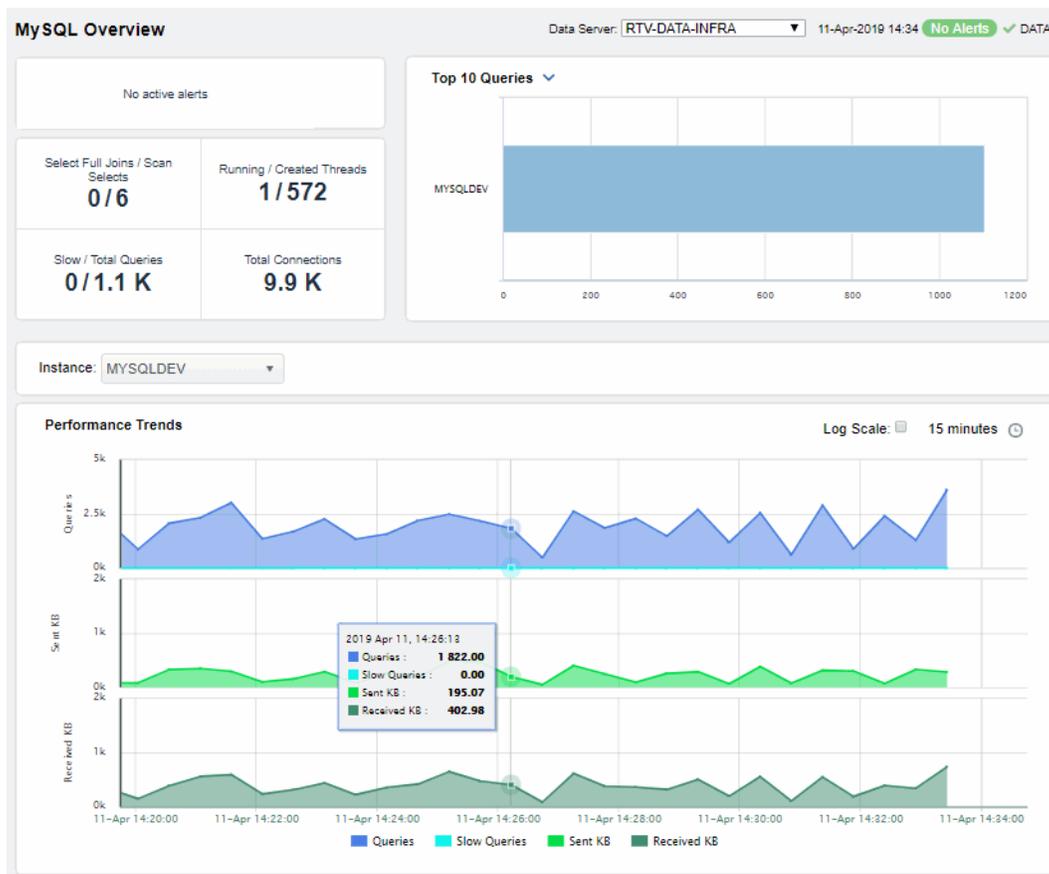
The total pending messages, the outgoing messages per second, and the incoming messages per second for a selected EMS Server on your connected DataServer.

You can hover over each region in the upper half of the Overview to see more detail. You can also drill down to see even more detail by clicking on each respective region in the Overview.

For example, clicking on the alerts in the CRITICAL and WARNING alerts region opens the Alerts Table display.

The bottom half of the display provides a performance trend graph for queries for a selected instance.

You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.



## All MySQL Instances View - HTML

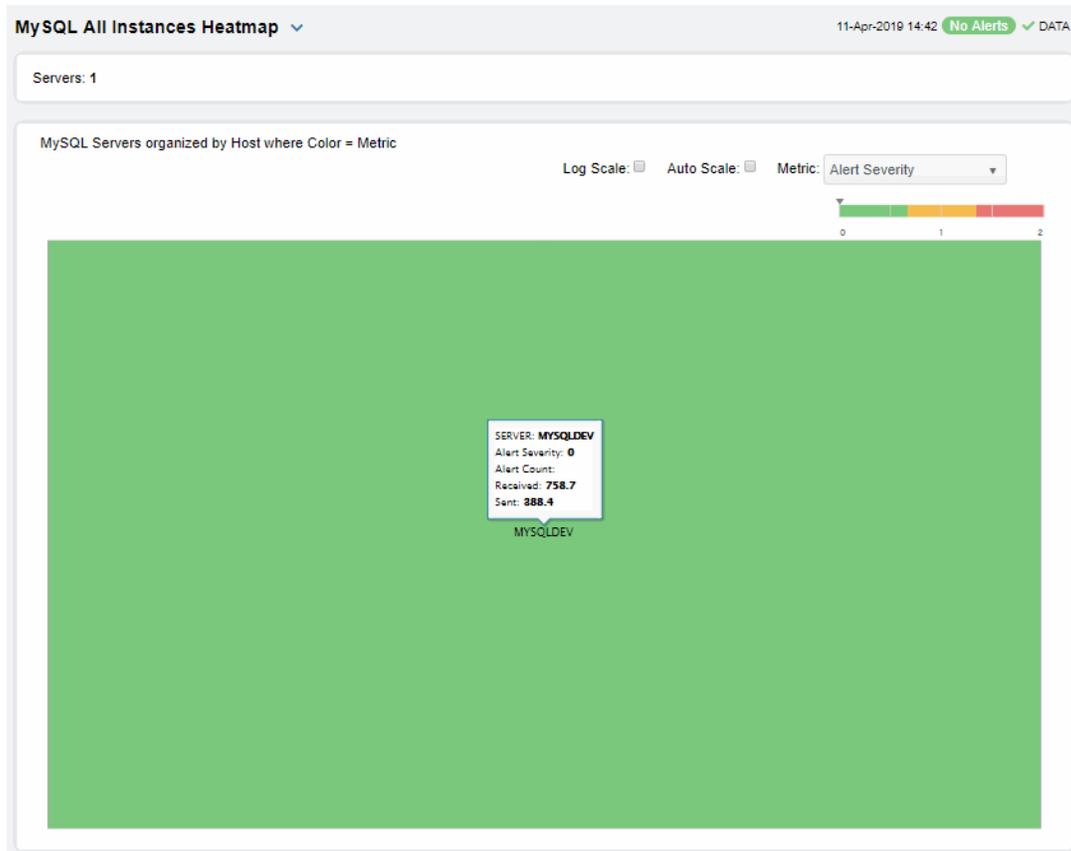
Displays in this View are:

- [“Instances Heatmap- HTML”](#): A heatmap view of all servers and their associated metrics.
- [“All MySQL Instances - HTML”](#): A tabular view of your servers and their associated metrics.

## Instances Heatmap- HTML

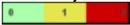
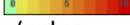
This heatmap display provides an easy-to-view interface that allows you to quickly identify the current status of each of your MySQL instances. Choose a metric from the **Metric** drop down menu. By default, this display shows the heatmap based on the **Alert Severity** metric. Other metrics are **Alert Count**, **Received** and **Sent**.

Each rectangle in the heatmap is a different MySQL instance. Mouse over a rectangle to see additional metrics for a server. Click a rectangle to open the ["Single MySQL Instance Summary- HTML"](#) display and see additional details for the selected MySQL instance.



### Fields and Data:

- Names** Select this check box to display the names of the instances at the top of each rectangle in the heatmap.
- Log** Select to this check box to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Auto Scale** Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value.  
**Note:** Some metrics auto-scale automatically, even when **Auto** is not selected.
- Metric** Choose a metric to view in the display. For details about the data, refer to vendor documentation.

<b>Alert Severity</b>	<p>The current alert severity. Values range from <b>0</b> - <b>2</b>, as indicated in the color gradient  bar, where <b>2</b> is the highest Alert Severity:</p> <ul style="list-style-type: none"> <li> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Alert Count</b>	<p>The total number of critical and warning unacknowledged alerts. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.</p>
<b>Received</b>	<p>The total number of bytes received. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the alarm threshold specified for the <b>MysqlBytesReceivedHigh</b> alert. The middle value in the gradient bar indicates the average count.</p>
<b>Sent</b>	<p>The total number of bytes sent. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the alarm threshold specified for the <b>MysqlBytesSentHigh</b> alert. The middle value in the gradient bar indicates the average count.</p>

## All MySQL Instances - HTML

Investigate detailed utilization metrics for all MySQL instances. This display provides a tabular view of the performance metrics shown in the [“Instances Heatmap- HTML”](#) (alert level, alert count, bytes received, and so forth), but with additional metrics such as **Delayed Writes**, **Queries**, **Connections**, **Time Stamp** and **Uptime**. Each row in the table contains data for a particular instance. Click a column header to sort column data in ascending or descending order. Double-click on a table row to drill-down to the [“Single MySQL Instance Summary- HTML”](#) display and view metrics for that particular instance. Toggle between the commonly accessed Table and Heatmap displays by clicking the drop down list on the display title.

Server	Expired	Alert Level	Alert Count	Bytes Received	Bytes Sent	Connections	Delayed Writes	Queries	Slow Queries
MYSQLDEV		✓		332,013	342,275	0,897	0	1,420	

## Single MySQL Instance View - HTML

Displays in this View are:

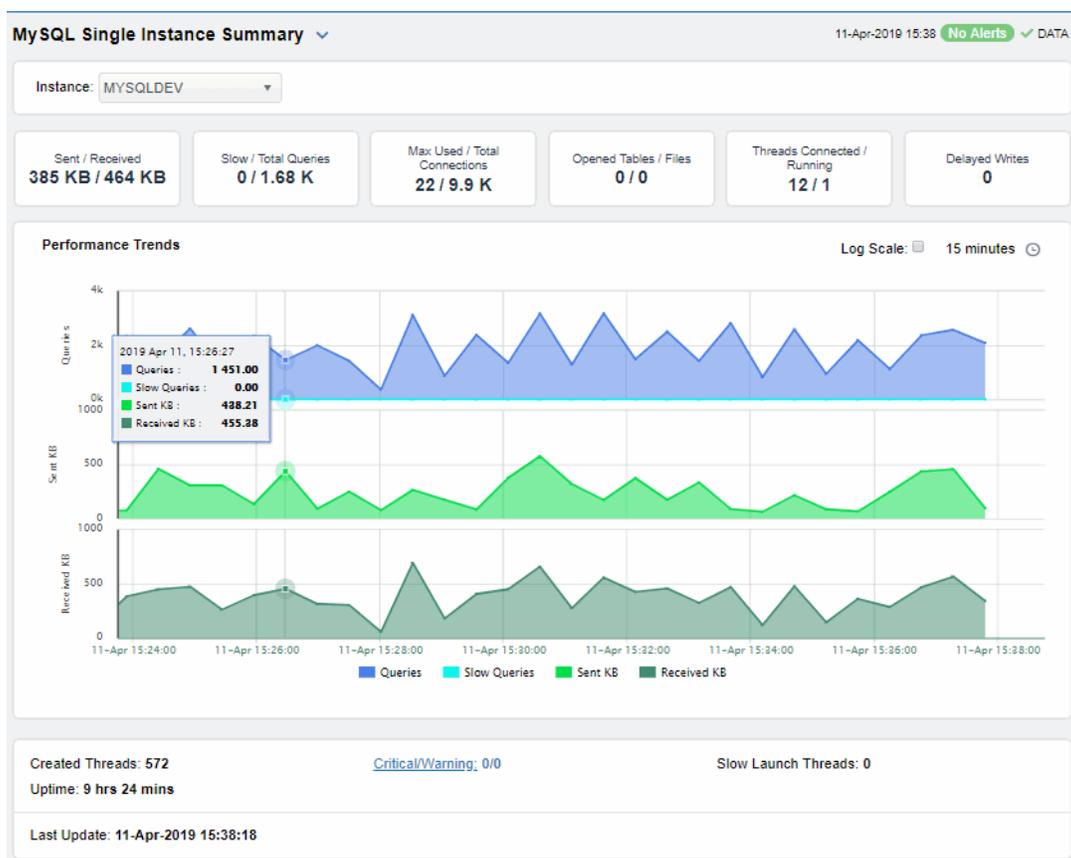
- [“Single MySQL Instance Summary- HTML”](#): Displays performance, processing, alerts, memory, and trend data for a particular database server.
- [“Instance Properties- HTML”](#): Displays the values of properties on servers.
- [“Instance Operations- HTML”](#): Trend graph that traces server queries, slow queries, KB sent and KB received.
- [“Instance User Tables- HTML”](#): A tabular view of cache tables performance and utilization metrics.

## Single MySQL Instance Summary- HTML

View connection, performance and processing details for a single MySQL instance, such as the total number of kilobytes sent and received, slow and total queries, maximum memory used, number of connections, opened tables and files, as well as the number of threads connected and running and delayed writes.

Choose an instance from the **Instance** drop-down menu. You can also drill up to see all MySQL instances in the ["All MySQL Instances - HTML"](#) display by clicking on values in the upper area.

The bottom half of the display provides a message rates trend graph for a selected MySQL instance. You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.



## Instance Properties- HTML

View properties and property values for a single MySQL instance.

Choose an instance from the **Instance** drop-down menu. Each table row is a different property for the selected instance. Enter a search string in the **Property Filter** field to limit the number of table rows. Click a column header to sort column data in numerical or alphabetical order.

MySQL Single Instance Properties 15-Apr-2019 08:48 Alerts DATA

Instance:

Filter by Property Name:

Property	Value
auto_increment_increment	1
auto_increment_offset	1
autocommit	ON
automatic_sp_privileges	ON
back_log	50
basedir	C:/Program Files/MySQL/MySQL Server 5.5/
big_tables	OFF
binlog_cache_size	32768
binlog_direct_non_transactional_updates	OFF
binlog_format	STATEMENT
binlog_stmt_cache_size	32768
bulk_insert_buffer_size	8388608
character_set_client	latin1
character_set_connection	latin1
character_set_database	latin1
character_set_filesystem	binary
character_set_results	latin1
character_set_server	latin1
character_set_system	utf8
character_sets_dir	C:/Program Files/MySQL/MySQL Server 5.5/share/charsets/
collation_connection	latin1_swedish_ci
collation_database	latin1_swedish_ci
collation_server	latin1_swedish_ci
completion_type	NO_CHAIN

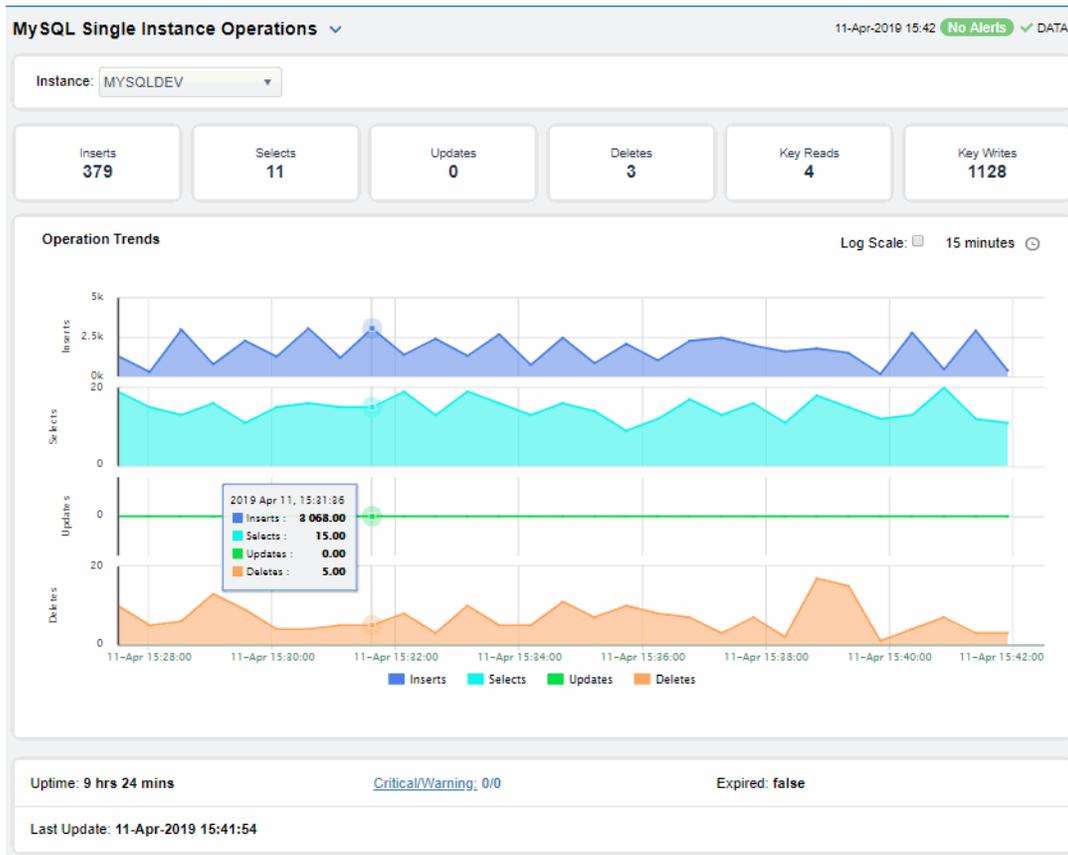
Page 1 of 7 1 - 40 of 259 items

## Instance Operations- HTML

View details about operations performed for a single MySQL instance, such as the total number of inserts, selects, updates, deletes, key reads and key writes. Choose an instance from the **Instance** drop-down menu. Click one of the values in the upper region to drill up to the ["All MySQL Instances - HTML"](#) display.

View trending performance data for a single MySQL instance: **Inserts**, **Selects**, **Updates** and **Deletes**. Choose an instance from the **Instance** drop-down menu. Mouse over the trend

graph to see performance metrics with time stamps.



### Instance User Tables- HTML

Investigate detailed utilization metrics for user tables on a single MySQL instance, such as **Data Size**, **Index Size**, **Row Count** and **Data Free**.

Each row in the table contains data for a particular user table on the selected MySQL instance. Click a column header to sort column data in ascending or descending order.

MySQL Single Instance User Tables 15-Apr-2019 08:49 No Alerts ✓ DATA

Instance: MYSQLDEV

Table	Data Size	Index Size	Row Count	Data Free	Total Size	Engine	Schema
alertlevels	0	1,024	0	0	1,024	MyISAM	alertdefs
audit_table	0	1,024	0	0	1,024	MyISAM	alertdefs
\$bw0_activities_table	47,221,756	9,963,488	515,918	0	56,885,244	MyISAM	rtvhistory
\$bw0_activity_totals_table	6,107,932	1,070,080	56,463	0	7,178,012	MyISAM	rtvhistory
\$bw0_process_totals_app_table	888,684	197,632	9,959	0	886,316	MyISAM	rtvhistory
\$bw0_process_totals_appnode_t	4,148,320	1,158,144	59,718	0	5,306,464	MyISAM	rtvhistory
\$bw0_process_totals_appslice_t	752,816	182,272	9,462	0	935,088	MyISAM	rtvhistory
\$bw0_process_totals_table	8,064,304	2,120,704	109,461	0	10,185,008	MyISAM	rtvhistory
\$bw0_processes_table	8,588,004	1,959,936	104,214	0	10,545,940	MyISAM	rtvhistory
bw0_activity_totals	20,718,016	4,355,072	226,128	0	25,073,088	MyISAM	rtvhistory
bw0_appnodes	2,597,056	764,928	39,409	0	3,361,984	MyISAM	rtvhistory
bw0_process_totals	7,650,588	1,859,584	94,395	0	9,510,172	MyISAM	rtvhistory
bw0_process_totals_app	777,800	216,064	10,979	0	993,864	MyISAM	rtvhistory
bw0_process_totals_appnode	4,415,924	1,270,784	65,919	0	5,686,708	MyISAM	rtvhistory
bw0_process_totals_appslice	5,211,584	1,274,880	65,961	0	6,486,464	MyISAM	rtvhistory
bw0_processes	0	2,048	0	0	2,048	MyISAM	rtvhistory
bw_activities	330,112,152	35,879,936	3,520,325	0	365,992,088	MyISAM	rtvhistory
bw_activity_totals	163,245,908	39,405,568	388,084	113,753,196	202,651,478	MyISAM	rtvhistory
bw_engines	15,505,928	4,278,272	35,389	11,208,024	19,784,200	MyISAM	rtvhistory
bw_process_totals	16,026,804	4,387,840	35,308	11,646,896	20,414,644	MyISAM	rtvhistory
bw_processes	202,308,492	40,863,744	389,137	145,412,016	243,172,236	MyISAM	rtvhistory
bw_servers	2,333,116	1,253,376	11,780	1,568,394	3,586,492	MyISAM	rtvhistory

Page 1 of 2 1 - 40 of 72 items

Tables: 72 Largest Table: **ems\_queuesext** Uptime Since Flush: 9 hrs 29 mins

## MS SQL Server

The Solution Package for Microsoft® SQL Server® includes high level heatmap and tabular displays as well as drilldown views to access real-time and historical performance metrics for each Microsoft SQL Server in your monitored services and applications.

With the Solution Package for Microsoft® SQL Server®, you are able to drill down from a high level alert at a business service or application health level into the supporting database infrastructure, to determine what is causing the alert and to take corrective action. This service-centric approach makes it easy for application support teams and Microsoft DBAs to prioritize incidents based on the impact to the business.

The following Views can be found under **Components** tab > **Databases** > **MS SQL Server**:

- **"All Servers View"**: The displays in this View allow you to view the current and historical metrics for all servers in a heatmap or tabular format.
- **"Single SQL Server View"**: The displays in this View allow you to view the metrics for a particular SQL database server.

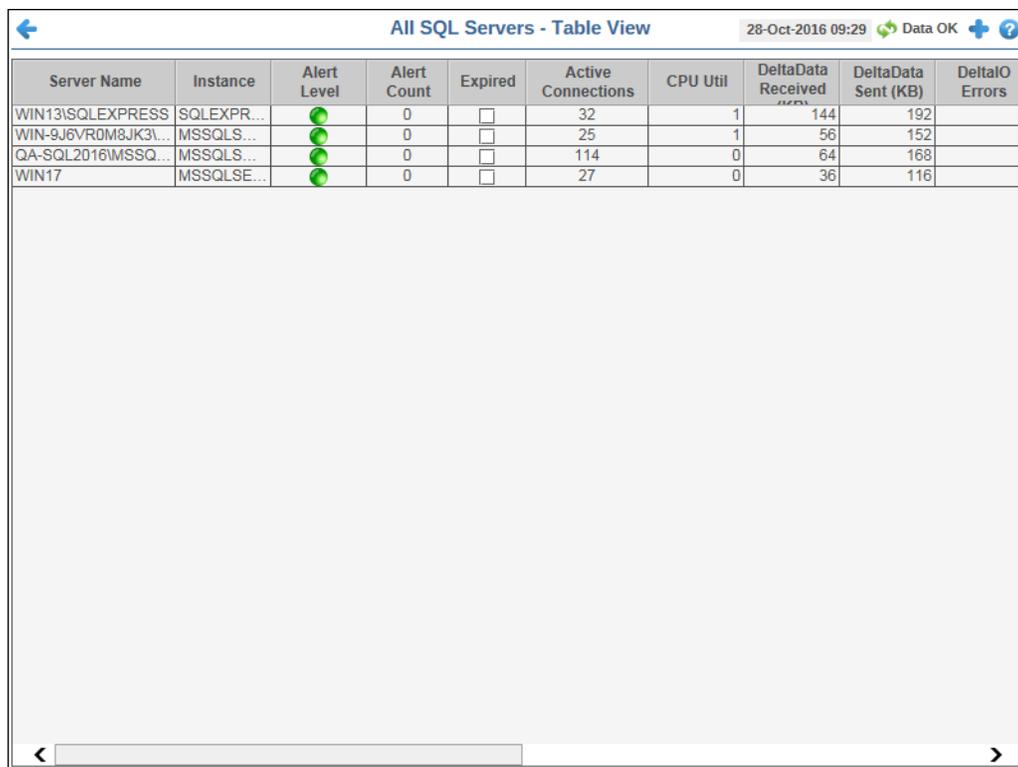
## All Servers View

These displays provide detailed data for all servers in a heatmap and tabular view. Displays in this View are:

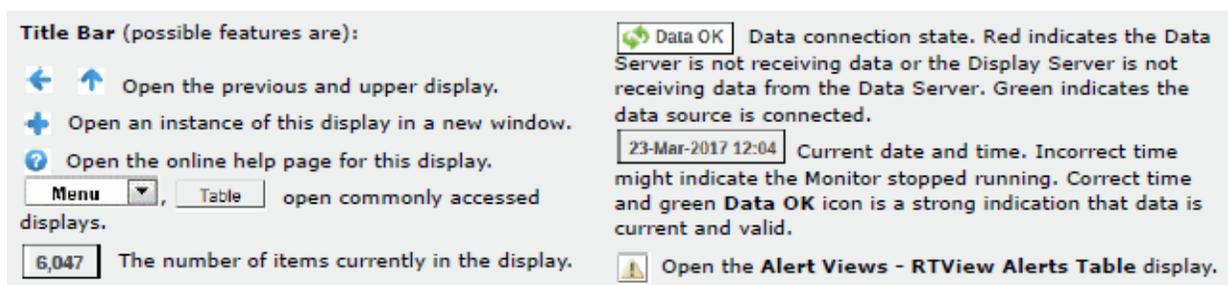
- ["All Servers"](#): A tabular view of your servers and their associated metrics.
- ["All Servers Heatmap"](#): A heatmap view of all servers and their associated metrics.

### All Servers

This table provides a view of all of your servers and their associated metric data including instance, alert severity, alert count, and the current value of each gathered metric. You can click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by clicking a row to view details for the selected server in the ["Server Summary"](#) display



Server Name	Instance	Alert Level	Alert Count	Expired	Active Connections	CPU Util	DeltaData Received	DeltaData Sent (KB)	DeltaIO Errors
WIN13\SQLEXPRESS	SQLEXPRESS	<span style="color: green;">●</span>	0	<input type="checkbox"/>	32	1	144	192	0
WIN-9J6VR0M8JK3\...	MSSQLS...	<span style="color: green;">●</span>	0	<input type="checkbox"/>	25	1	56	152	0
QA-SQL2016\MSSQ...	MSSQLS...	<span style="color: green;">●</span>	0	<input type="checkbox"/>	114	0	64	168	0
WIN17	MSSQLSE...	<span style="color: green;">●</span>	0	<input type="checkbox"/>	27	0	36	116	0



**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the server. Refer to Microsoft SQL Server documentation for more information regarding these fields.

## All SQL Servers Table

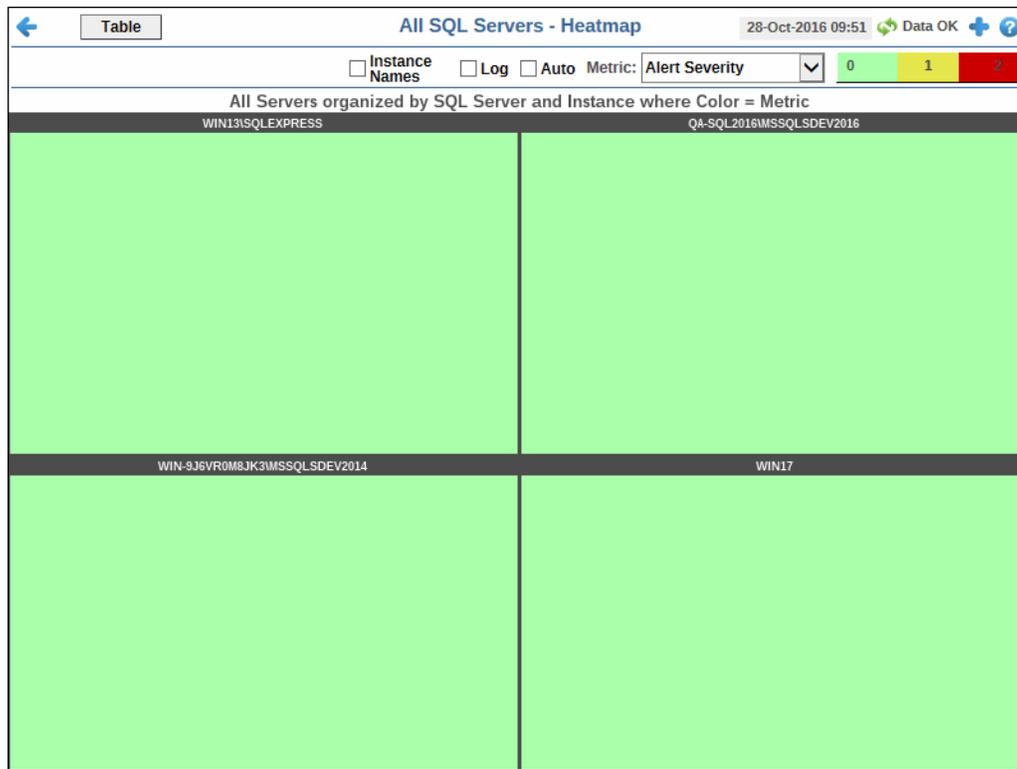
<b>Server Name</b>	The name of the server.
<b>Instance</b>	The name of the instance.
<b>Alert Level</b>	The current alert severity. <ul style="list-style-type: none"> <li><span style="color: red;">●</span> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li><span style="color: yellow;">●</span> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li><span style="color: green;">●</span> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Alert Count</b>	The total number of alerts for the host.
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > ( <b>Project Name</b> ) > <b>Solution Package Configuration</b> > <b>Microsoft SQL Server</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Active Connections</b>	The number of currently active connections.
<b>CPU Util</b>	The CPU utilization percentage.*
<b>Delta Data Received (KB)</b>	The increase in the amount of data being received (from the previous polling period to the current polling period), in kilobytes.
<b>Delta Data Sent (KB)</b>	The increase in the amount of data being sent (from the previous polling period to the current polling period), in kilobytes.
<b>Delta IO Errors</b>	The increase in the amount of input/output errors (from the previous polling period to the current polling period).
<b>Delta IO Reads</b>	The increase in the amount of input/output reads operations (from the previous polling period to the current polling period).
<b>Delta IO Writes</b>	The increase in the amount of input/output write operations (from the previous polling period to the current polling period).
<b>Delta Packet Errors</b>	The increase in the amount of packet errors (from the previous polling period to the current polling period).

<b>IO Busy (ms)</b>	The time, in milliseconds, that the system has been busy due to Input/Output operations.*
<b>Memory Used (%)</b>	The percentage of memory used on the server.*
<b>Memory Remaining (%)</b>	The percentage of memory remaining on the server.*
<b>Total DB Size (MB)</b>	The size of the database, in megabytes.*
<b>Server Edition</b>	The version of the server.*
<b>Product Level</b>	The product level of the server.*
<b>Product Version</b>	The product's version number.*

## All Servers Heatmap

This heatmap display provides an easy-to-view interface that allows you to quickly identify the current status of each of your servers for each available metric. You can view the servers in the heatmap based on the following metrics: the current alert severity, the current alert count, and the percentage of CPU used. By default, this display shows the heatmap based on the **Alert Severity** metric.

You can use the **Instance Names** check-box  to include or exclude labels in the heatmap, and you can mouse over a rectangle to see additional metrics for an engine. Clicking one of the rectangles in the heatmap opens the "[Server Summary](#)" display, which allows you to see additional details for the selected server.



#### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.

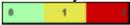
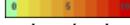
**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

Open the **Alert Views - RTView Alerts Table** display.

#### Fields and Data:

- Instance Names** Select this check box to display the names of the instances at the top of each rectangle in the heatmap.
- Log** Select this check box to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Auto** Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value.  
**Note:** Some metrics auto-scale automatically, even when **Auto** is not selected.
- Metric** Choose a metric to view in the display.

<b>Alert Severity</b>	<p>The current alert severity. Values range from <b>0</b> - <b>2</b>, as indicated in the color gradient  bar, where <b>2</b> is the highest Alert Severity:</p> <ul style="list-style-type: none"> <li> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Alert Count</b>	<p>The total number of critical and warning unacknowledged alerts in the engine. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.</p>
<b>SQL CPU Utilization (%)</b>	<p>The percentage of CPU used by the instance. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of <b>MssqlInstanceSqlCpuUsageHigh</b>. The middle value in the gradient bar indicates the middle value of the range.</p> <p>When <b>Auto</b> is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>

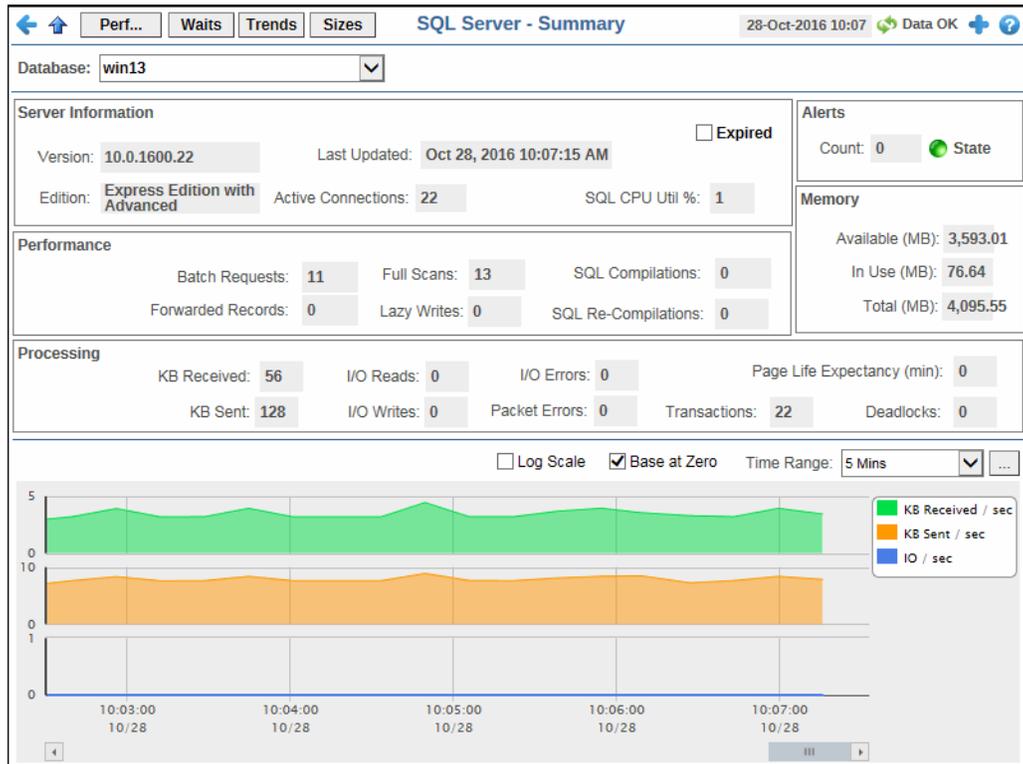
## Single SQL Server View

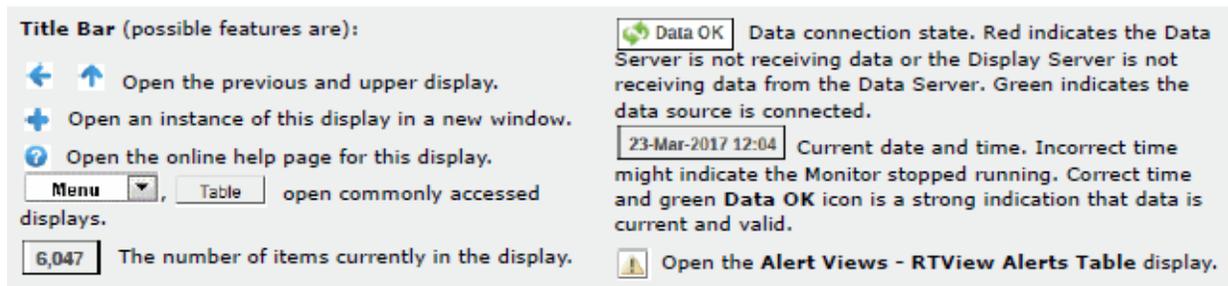
Displays in this View are:

- **"Server Summary"**: Displays performance, processing, alerts, memory, and trend data for a particular database server.
- **"Database Details"**: Displays various database details as well as trending data for the page life expectancy.
- **"Wait Stats"**: Displays server wait time details in a table format for a particular database server.
- **"Performance Trends"**: This display allows you to view performance trend data for a particular SQL database server.
- **"Database/Table Sizes"**: Displays database and table sizes for a particular database server.

## Server Summary

This display allows you to view connection and CPU utilization details, memory statistics, various performance and processing metrics, and trending data for the number of kilobytes received and sent as well as input/output details per second for a particular SQL database server.






---

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the selected server. Refer to Microsoft SQL Server documentation for more information regarding these fields.

---



---

**Note:** The **Perf** button takes you to "Database Details". The **Waits** button takes you to "Wait Stats". The **Trends** button takes you to "Performance Trends". The **Sizes** button takes you to "Database/Table Sizes".

---

#### Filter By:

**Database** Select the database for which you want to show data in the display.

#### Fields and Data:

##### Server Information

<b>Version</b>	The server's version number
<b>Last Updated</b>	The date and time of the last data update.
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > ( <b>Project Name</b> ) > <b>Solution Package Configuration</b> > <b>Microsoft SQL Server</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Edition</b>	The SQL Server's edition.*
<b>Active Connections</b>	The number of active connections on the server.*
<b>SQL CPU Util %</b>	The percentage of CPU used by the server.*

##### Alerts

**Count** The total number of current alerts.

<b>State</b>	The current alert severity.  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
--------------	---

**Memory**

<b>Available (MB)</b>	The amount of memory currently available, in megabytes.*
<b>In Use (MB)</b>	The amount of memory currently in use, in megabytes.*
<b>Total (MB)</b>	The total amount of memory, in megabytes.*

**Processing**

<b>KB Received</b>	The number of kilobytes received.*
<b>KB Sent</b>	The number of kilobytes sent.*
<b>I/O Reads</b>	The number of input/output reads.*
<b>I/O Writes</b>	The number of input/output writes.*
<b>I/O Errors</b>	The number of input/output errors.*
<b>Packet Errors</b>	The number of errors involving incoming/outgoing packets.*
<b>Page Life Expectancy (min)</b>	The average number of minutes a page stays in the cache.*
<b>Transactions</b>	The number of transactions.*
<b>Deadlocks</b>	The number of deadlocks.*

**Performance Trends Graph**

Traces the following:

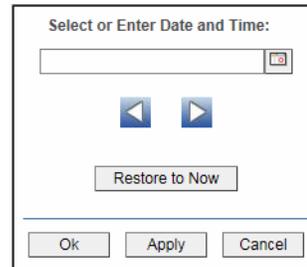
**KB Received/sec** -- traces the amount of kilobytes received per second.

**KB Sent/sec** -- traces the amount of kilobytes sent per second.

**I/O /sec**-- traces the average number of input/output operations per second.

**Log Scale** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.

- Base at Zero** Select to use zero (**0**) as the Y axis minimum for all graph traces.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



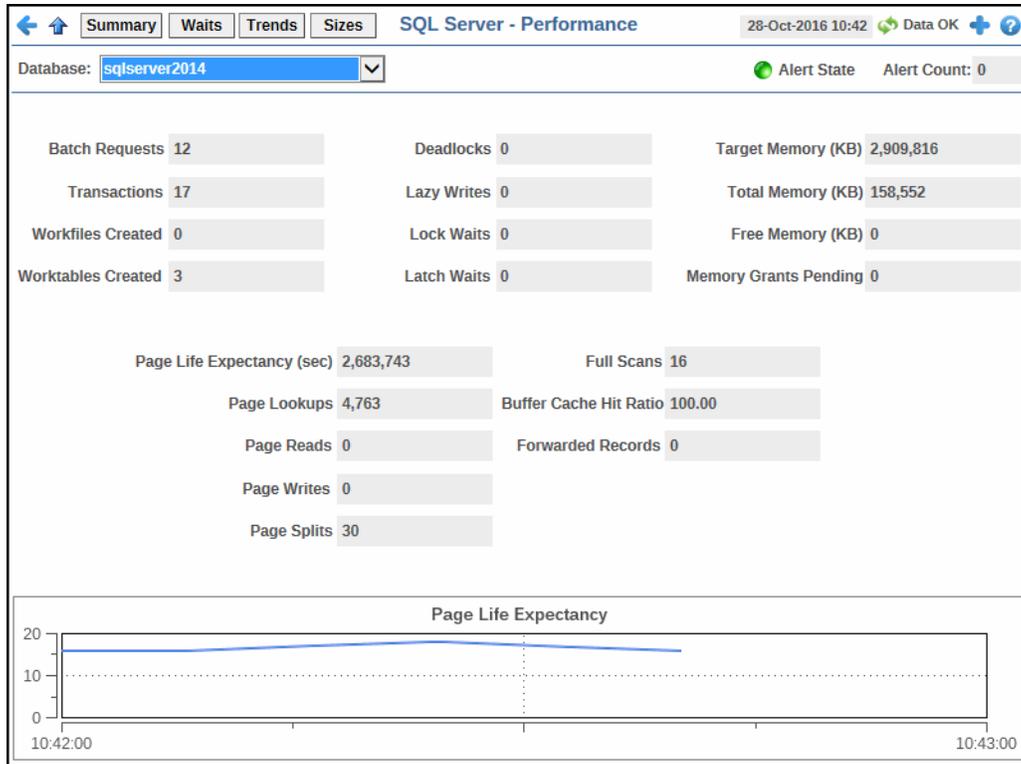
By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Database Details

This display allows you to view various database details as well as trending data for the page life expectancy.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

Open the **Alert Views - RTView Alerts Table** display.

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the selected server. Refer to Microsoft SQL Server documentation for more information regarding these fields.

**Note:** The **Summary** button takes you to "Server Summary". The **Waits** button takes you to "Wait Stats". The **Trends** button takes you to "Performance Trends". The **Sizes** button takes you to "Database/Table Sizes".

**Filter By:**

**Database** Select the database for which you want to show data in the display.

**Fields and Data:**

<b>Alert State</b>	The current alert severity.  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
<b>Alert Count</b>	The total number of current alerts.
<b>Batch Requests</b>	The current number of batch requests.*
<b>Transactions</b>	The current number of transactions.*
<b>Workfiles Created</b>	The number of work files created.*
<b>Worktables Created</b>	The number of worktables created.*
<b>Deadlocks</b>	The current number of deadlocks occurring in the database.*
<b>Lazy Writes</b>	The number of times per second SQL Server relocates dirty pages from buffer pool (memory) to disk.*
<b>Lock Waits</b>	The number of lock requests that required the caller to wait.*
<b>Latch Waits</b>	The number of latch requests that required the caller to wait.*
<b>Target Memory (KB)</b>	The defined target server memory, which is the ideal amount of memory the server can consume, in kilobytes.*
<b>Total Memory (KB)</b>	The total amount of memory the server has committed using the memory manager, in kilobytes.*
<b>Free Memory (KB)</b>	The total amount of free memory, in kilobytes.*
<b>Memory Grants Pending</b>	The current number of processes waiting for a workspace memory grant.*
<b>Page Life Expectancy (sec)</b>	The average number of seconds a page stays in the cache.*
<b>Page Lookups</b>	The number of page lookups.*
<b>Page Reads</b>	The number of database pages read per second.*
<b>Page Writes</b>	The number of database pages written per second.*
<b>Page Splits</b>	The number of page splits that occur as a result of overflowing index pages.*
<b>Full Scans</b>	The number of full database scans.*
<b>Buffer Cache Hit Ratio</b>	The current buffer cache hit ratio, which is the total number of cache hits divided by the total number of cache lookups.*

**Forwarded Records**

The number of records fetched through forward record pointers.\*

**Page Life Expectancy Trend Graph**

Traces the average length of time a page stays in the cache.\*

**Wait Stats**

This display allows you to view server wait time details in a table format for a particular database server. You can drill-down and view the details for a particular container in the "Server Summary" display by clicking on a row in the resulting table.

Database: sqlserver2014      Alert State      Alert Count: 0

Wait Types Sorted By Non-Zero Percentage Per Wait Category

Wait Category	Wait Type	Wait (sec)	Signal (sec)	Resource (s...	Percentage	Wait Count	Avg Wait (sec)
I/O		0.05	0.00	0.05	0.01	19.00	0.04
Lock		0.04	0.00	0.04	0.00	2.00	0.03
Non-I/O Page La...		0.00	0.00	0.00	0.00	115.00	0.04
Other		0.00	0.00	0.00	0.00	2.00	0.04
Transaction Log		0.00	0.00	0.00	0.00	2.00	0.04

**Title Bar** (possible features are):

-   Open the previous and upper display.
-  Open an instance of this display in a new window.
-  Open the online help page for this display.
- open commonly accessed displays.
- The number of items currently in the display.

 **Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

 Open the **Alert Views - RTView Alerts Table** display.

---

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the selected server. Refer to Microsoft SQL Server documentation for more information regarding these fields.

---



---

**Note:** The **Summary** button takes you to "[Server Summary](#)". The **Perfs** button takes you to "[Database Details](#)". The **Trends** button takes you to "[Performance Trends](#)". The **Sizes** button takes you to "[Database/Table Sizes](#)".

---

#### Filter By:

The display includes these filtering options:

<b>Database</b>	Select the database for which you want to show data in the display.
<b>Alert State</b>	The current alert status. <ul style="list-style-type: none"> <li> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Alert Count</b>	Total number of alerts for the process.

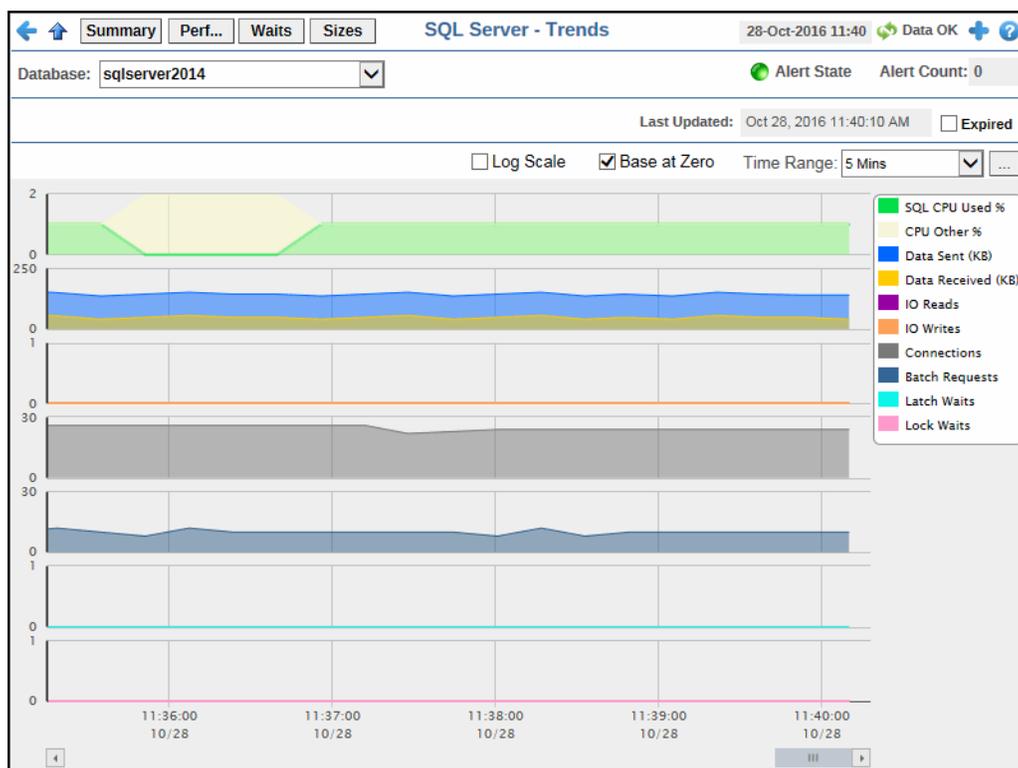
#### Wait Types Sorted By Non-Zero Percentage Per Wait Category Table

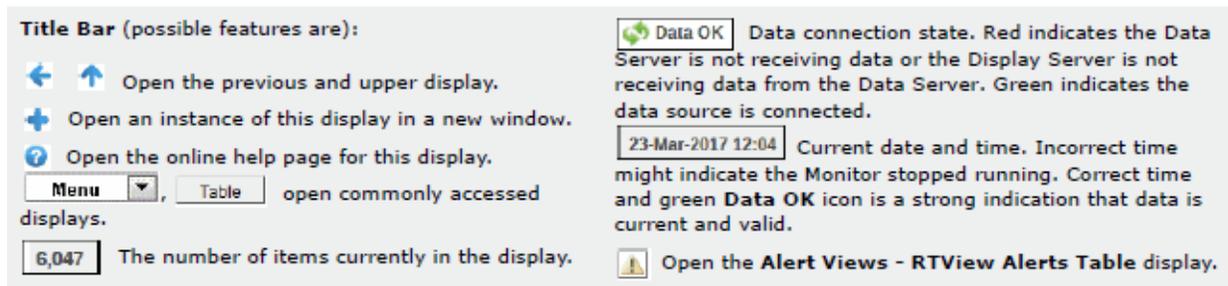
<b>Wait Category</b>	The name of the wait category.*
<b>Wait Type</b>	The name of the wait type.*
<b>Wait (sec)</b>	The average length of the wait time, in seconds.*
<b>Signal (sec)</b>	When the thread is marked as runnable, this field displays the wait time, in seconds, that it takes to get into the running state.*
<b>Resource (sec)</b>	The length of time the thread spent in a suspended state waiting to acquire a resource, in seconds.*
<b>Percentage</b>	The percentage of time the thread spent in a wait state for this wait type.*
<b>Wait Count</b>	The number of lock requests that required the caller to wait.*
<b>Avg Wait (sec)</b>	The average wait time, in seconds.*

<b>Avg Signal (sec)</b>	The average wait signal time, in seconds.*
<b>Avg Resource (sec)</b>	The average length of time taken to acquire a resource, in seconds.*
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > ( <b>Project Name</b> ) > <b>Solution Package Configuration</b> > <b>Microsoft SQL Server</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Timestamp</b>	The date and time the row data was last updated.

## Performance Trends

This display traces the current and historical percentage of CPU used by the MS SQL Server, the remainder CPU used in other operations, the amount of data sent, the amount of data received, the number of input/output operation reads, the number of input/output operation writes, the number of connections, the number of batch requests, the number of latch waits, and the number of lock waits for a particular SQL database server.






---

**Note:** The **Summary** button takes you to "Server Summary". The **Perf** button takes you to "Database Details". The **Waits** button takes you to "Wait Stats". The **Sizes** button takes you to "Database/Table Sizes".

---

### Filter By:

The display might include these filtering options:

- Database** Select the database for which you want to show data in the display.
- Alert State** The current alert status.
  - Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
  - Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
  - Green indicates that no metrics have exceeded their alert thresholds.
- Alert Count** Total number of alerts for the process.
- Last Updated** The date and time the data in the display was last updated.
- Expired** When checked, performance data has not been received within the time specified (in seconds) in the **Expire Time** field in the **Duration** region in the RTView Configuration Application > (**Project Name**) > **Solution Package Configuration** > **Microsoft SQL Server** > **DATA STORAGE** tab. The **Delete Time** field (also in the **Duration** region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
- Performance Trends Graph** Traces the following:
  - SQL CPU Used %** -- traces percentage of CPU used by the MS SQL Server.
  - CPU Other %** -- traces the percentage of CPU used in other operations by the MS SQL Server.
  - Data Sent (KB)** -- traces the amount of data sent, in kilobytes.
  - Data Received (KB)** -- traces the amount of data received, in kilobytes.
  - IO Reads** -- traces the number of input/output operation reads.
  - IO Writes** -- traces the number of input/output operation writes.
  - Connections** -- traces the number of connections.
  - Batch Requests** -- traces the number of batch requests.
  - Latch Waits** -- traces the number of latch waits.
  - Lock Waits** -- traces the number of lock waits.

**Log Scale**

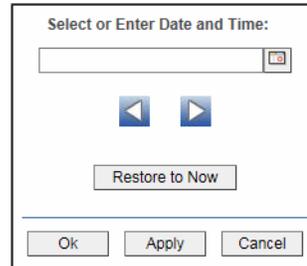
Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.

**Base at Zero**

Select to use zero (**0**) as the Y axis minimum for all graph traces.

**Time Range**

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Database/Table Sizes

This display provides database and table size data for a particular SQL database server.

**Database Sizes**

Database	Expired	Log Size (MB)	Row Size (MB)	State	Total Size (MB)
model	<input type="checkbox"/>	0.75	3.19	ONLINE	3.94
jparker	<input type="checkbox"/>	1	4	ONLINE	5
dcavasos	<input type="checkbox"/>	1	4	ONLINE	5
ReportServer\$MSSQL...	<input type="checkbox"/>	1.06	4.19	ONLINE	5.25
master	<input type="checkbox"/>	2	4	ONLINE	6
tempdb	<input type="checkbox"/>	0.5	8	ONLINE	8.5
ReportServer\$MSSQL...	<input type="checkbox"/>	7.56	5.19	ONLINE	12.75
msdb	<input type="checkbox"/>	19.63	17.13	ONLINE	36.75

**Table Sizes**

Database	Schema	Table	Rows	Total (MB)	Used (MB)	Unused (MB)	Creation Date
jparker	dbo	EMS_TOPICTOTALS	0	0	0	0	Sep 27, 2016 1...
jparker	dbo	EMS_TOPICSEXT	0	0	0	0	Sep 27, 2016 1...
jparker	dbo	EMS_TOPICS	0	0	0	0	Sep 27, 2016 1...
jparker	dbo	EMS_SERVERINFOEXT	0	0	0	0	Sep 27, 2016 1...
jparker	dbo	EMS_SERVERINFO	0	0	0	0	Sep 27, 2016 1...
jparker	dbo	EMS_ROUTES	0	0	0	0	Sep 27, 2016 1...
jparker	dbo	EMS_ROUTECOUNTS	0	0	0	0	Sep 27, 2016 1...
jparker	dbo	EMS_QUEUETOTALS	0	0	0	0	Sep 27, 2016 1...
jparker	dbo	EMS_QUEUESEXT	0	0	0	0	Sep 27, 2016 1...
jparker	dbo	EMS_QUEUEES	0	0	0	0	Sep 27, 2016 1...
jparker	dbo	EMS_PRODUCERS	0	0	0	0	Sep 27, 2016 1...
jparker	dbo	EMS_PRODUCERS*	0	0	0	0	Sep 27, 2016 1...

**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** , **Table** open commonly accessed displays.
- 6,047** The number of items currently in the display.
- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the selected server. Refer to Microsoft SQL Server documentation for more information regarding these fields.

**Note:** The **Summary** button takes you to "Server Summary". The **Perf** button takes you to "Database Details". The **Waits** button takes you to "Wait Stats". The **Trends** button takes you to "Performance Trends".

**Filter By:**

The display includes these filtering options:

<b>Database</b>	Select the database for which you want to show data in the display.
<b>Alert State</b>	The current alert status.  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
<b>Alert Count</b>	Total number of alerts for the process.
<b>Last Updated</b>	The date and time the data in the display was last updated.
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > ( <b>Project Name</b> ) > <b>Solution Package Configuration</b> > <b>Microsoft SQL Server</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.

**Database Sizes Table**

<b>Database</b>	The name of the database.*
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > ( <b>Project Name</b> ) > <b>Solution Package Configuration</b> > <b>Microsoft SQL Server</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Log Size (MB)</b>	The size of the log, in megabytes.*
<b>Row Size (MB)</b>	The row size, in megabytes.*
<b>State</b>	The current state of the database.*
<b>Total Size (MB)</b>	The total size of the database, in megabytes.*

**Table Sizes**

<b>Database</b>	The name of the database.*
<b>Schema</b>	The name of the schema.*
<b>Table</b>	The name of the table.*
<b>Rows</b>	The number of rows in the table.*
<b>Total (MB)</b>	The total, in megabytes, available in the table.*
<b>Used (MB)</b>	The total number of used megabytes in the table.*
<b>Unused (MB)</b>	The total number unused megabytes in the table.*
<b>Creation Date</b>	The date and time the table was created.*

---

## MS SQL - HTML

The following can be found under **Components** tab > **Databases** > **MS SQL Server**:

- [“MS SQL Overview - HTML”](#):
- [“MS Servers View - HTML”](#): The displays in this View allow you to view the current and historical metrics for all servers in a heatmap or tabular format.
- [“MS SQL Server View - HTML”](#): The displays in this View allow you to view the metrics for a particular SQL database server.

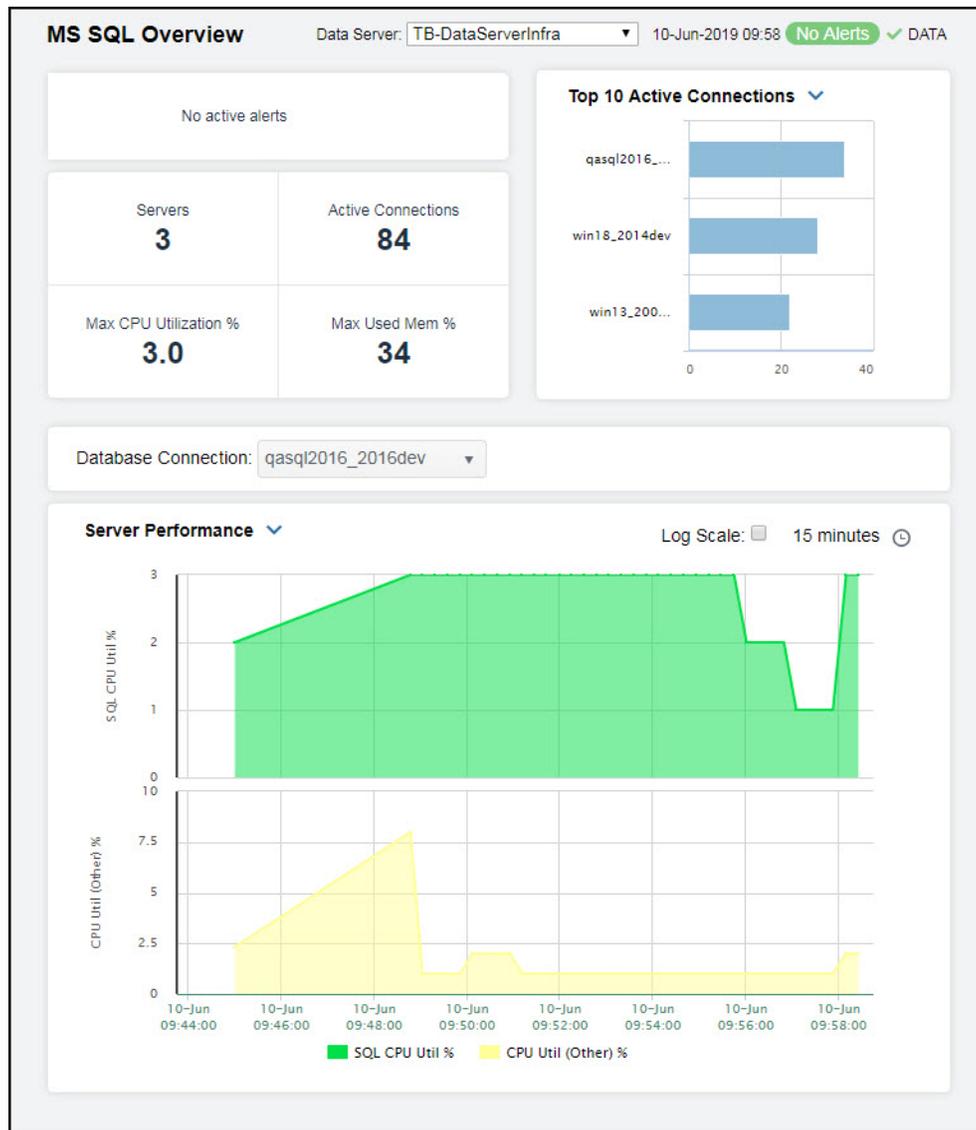
### MS SQL Overview - HTML

The **MS SQL Overview** is the top-level display for the MS SQL Monitor, which provides a good starting point for immediately getting the status of all your database instances and connections (on those instances) on your Data Server. You can select the RTView DataServer for which you want to see data and easily view the current data for that DataServer including:

- The total number of active alerts for the selected DataServer, including the total number of critical and warning alerts.
- The number of database instances and the number of connections on those instances.
- The maximum percentage of CPU utilization across all database instances.
- The maximum percentage of memory utilization across all database instances.
- A visual list of the top 10 active connections, connections with the best CPU utilization percentage, connections with the best used memory percentage, connections with the largest DB size, connections with the most received kilobytes, and connections with the most sent kilobytes on your connected DataServer.

You can hover over each region in the upper half of the Overview to see more detail. You can also drill down to see even more detail by clicking on each respective region in the Overview. For example, clicking on the alerts in the **CRITICAL** and **WARNING** alerts region opens the **Alerts Table by Components** display.

The bottom half of the display provides a host resources trend graph for a selected database connection. You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.



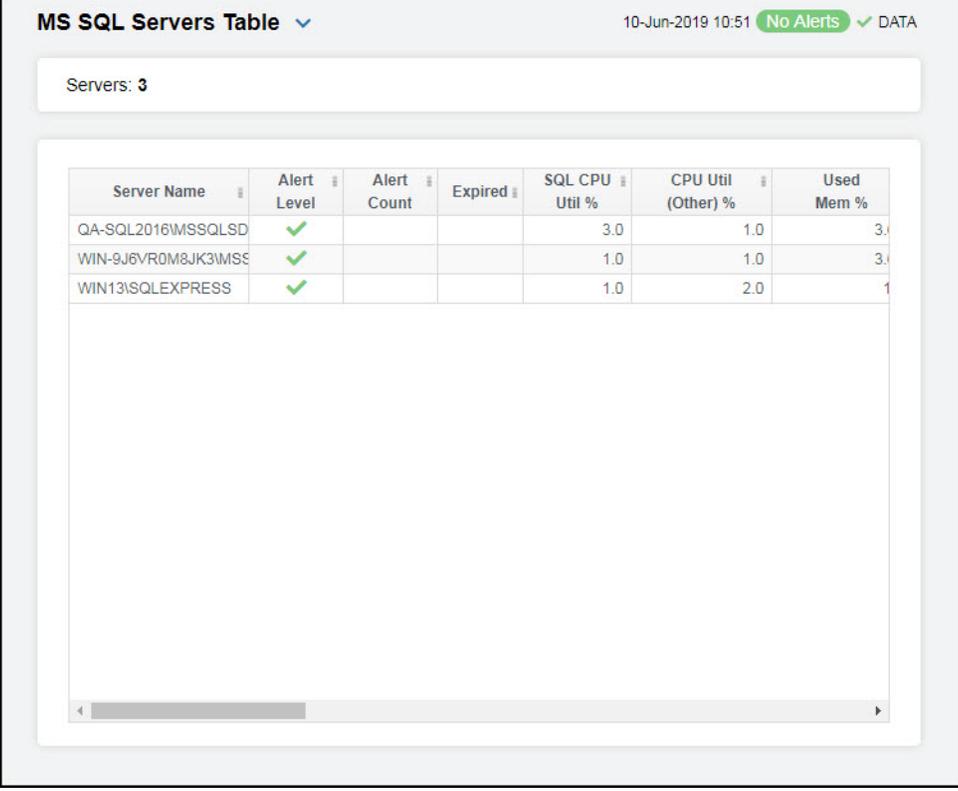
## MS Servers View - HTML

These displays provide detailed data for all MS SQL servers. Clicking MS SQL Servers from the left/navigation menu opens the ["MS SQL Servers Table - HTML"](#) display, which shows a tabular view of all MS SQL servers and their associated metrics. The option available under **MS SQL Servers** is:

- **MS SQL Servers Heatmap**: Opens the ["MS SQL Servers Heatmap - HTML"](#) display, which provides a heatmap view of all MS SQL servers and their associated metrics.

## MS SQL Servers Table - HTML

This table provides a view of all of your servers and their associated metric data including instance, alert severity, alert count, and the current value of each gathered metric. You can click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by double-clicking a row to view details for the selected server in the ["MS SQL Server Summary - HTML"](#) display.



MS SQL Servers Table 10-Jun-2019 10:51 No Alerts DATA

Servers: 3

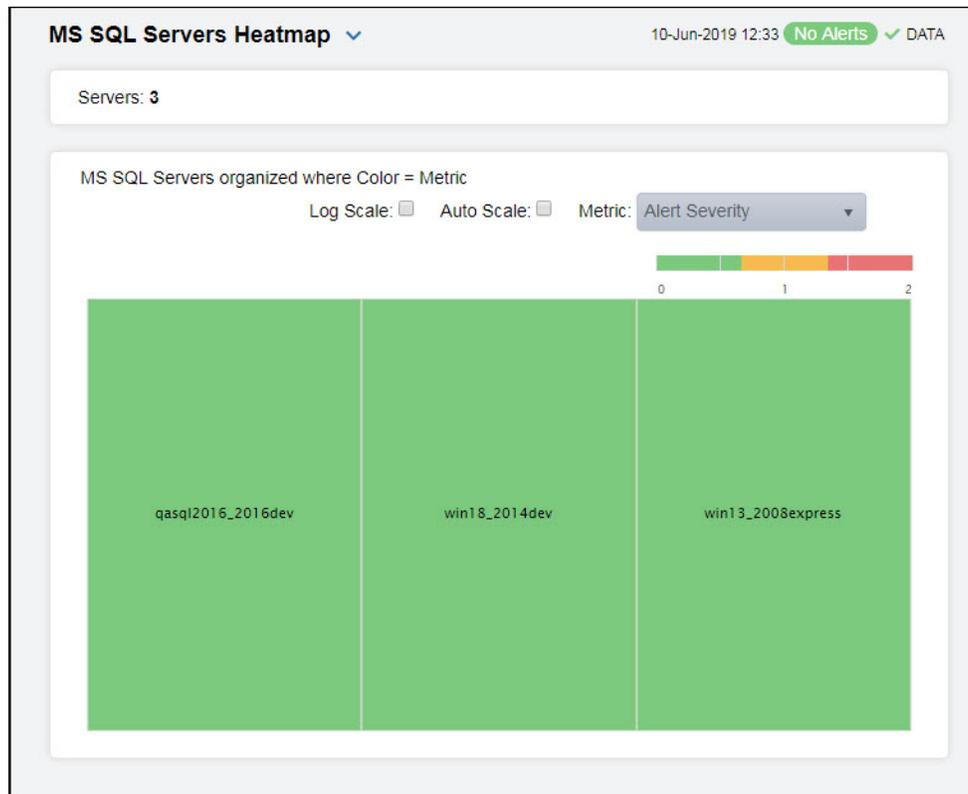
Server Name	Alert Level	Alert Count	Expired	SQL CPU Util %	CPU Util (Other) %	Used Mem %
QA-SQL2016WSSQLSD	OK			3.0	1.0	3.0
WIN-9J6VR0M8JK3MSE	OK			1.0	1.0	3.0
WIN13\SQLEXPRESS	OK			1.0	2.0	1.0

## MS SQL Servers Heatmap - HTML

Clicking **MS SQL Servers Heatmap** in the left/navigation menu opens the **MS SQL Servers Heatmap**, which provides an easy-to-view interface that allows you to quickly identify the current status of each of your servers for each available metric. You can view the servers in the heatmap based on the following metrics: the current alert severity, the current alert count, and the percentage of CPU used. By default, this display shows the heatmap based on the **Alert Severity** metric.

Each rectangle in the heatmap represents a server. The rectangle color indicates the most critical alert state associated with the server. Choose a different metric to display from the **Metric** drop-down menu. Mouse over a rectangle to see additional metrics. By default, this display shows **Alert Severity**.

Drill-down and investigate a server by clicking a rectangle in the heatmap to view details in the ["MS SQL Server Summary - HTML"](#) display.



### Available Metrics

Select the metric driving the heatmap display. The default is **Alert Severity**. Each **Metric** has a color gradient bar that maps values to colors. The heatmap organizes the servers by host, where each rectangle represents an server. Mouse-over any rectangle to display the current values of the metrics for the server. Click on a rectangle to drill-down to the associated ["MS SQL Server Summary - HTML"](#) display for a detailed view of metrics for that particular server.

**Alert Severity** The current alert severity. Values range from **0 - 2**, as indicated in the color gradient bar, where **2** is the highest Alert Severity:

- Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
- Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
- Green indicates that no metrics have exceeded their alert thresholds.

**Alert Count** The total number of critical and warning unacknowledged alerts in the engine. The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.

**SQL CPU Utilization (%)** The percentage of CPU used by the instance. The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **MssqlInstanceSqlCpuUsageHigh**. The middle value in the gradient bar indicates the middle value of the range.

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

## MS SQL Server View - HTML

These displays provide detailed data for a particular server. Clicking **MS SQL Server** from the left/navigation menu opens the "[MS SQL Server Summary - HTML](#)" display, which displays performance, processing, alerts, memory, and trend data for a particular database server. The options available under **MS SQL Server** are:

- **Instance Details:** Opens the "[MS SQL Server Performance - HTML](#)" display, which displays various database details as well as trending data for the page life expectancy.
- **Wait Stats:** Opens the "[MS SQL Server Waits Table - HTML](#)" display, which displays server wait time details in a table format for a particular database server.
- **Table Sizes:** Opens the "[MS SQL Server DB Table Sizes - HTML](#)" display, which displays database and table sizes for a particular database server.

## MS SQL Server Summary - HTML

Clicking **MS SQL Server** in the left/navigation menu opens the **MS SQL Server Summary** display, which allows you to view current as well as trending data for the server. Clicking on the information boxes at the top of the display takes you to the "[MS SQL Servers Table - HTML](#)" display, where you can view additional engines data.

The trend graph at the bottom of the display contains three different options: **Server Performance**, **Server Throughput**, and **Server Operations**. The **Server Performance** trend graph allows you to view trend data for the SQL CPU utilization percentage and CPU utilization (other types) over a selected time range. The **Server Throughput** trend graph allows you to view trend data for current data being sent/received and the current input/output reads/writes over a selected time range. The **Server Operations** trend graph allows you to view trend data for active connections, current batch requests, current latch waits, and current lock waits over a selected time range.

Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display.

**MS SQL Server Summary** ▾
10-Jun-2019 14:31 No Alerts DATA

Database Connection: gasql2016\_2016dev ▾

Active Connections  
**28**

SQL CPU Util %  
**3.0**

Memory Used %  
**3.7**

Data Sent KB / Received KB  
**344.0 / 84.0**

Total DB Size MB  
**25,483.3**

Deadlocks  
**0**

**Server Performance** ▾
Log Scale:  15 minutes ⌵

■ SQL CPU Util %
■ CPU Util (Other) %

Server Name: **QA-SQL2016\MSSQLSDEV2016** Instance: **MSSQLSDEV2016**

---

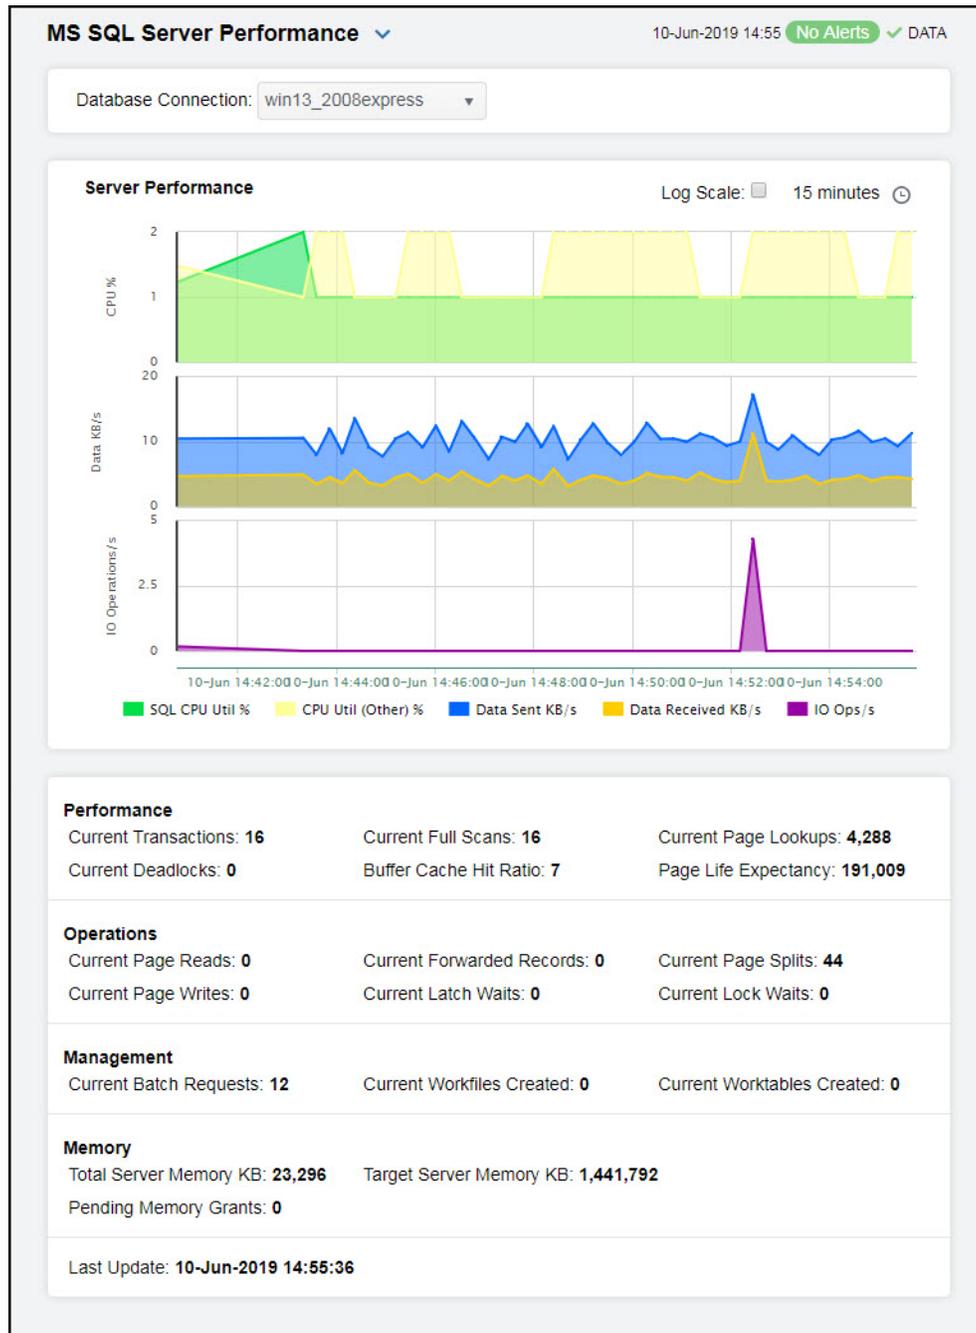
Product Version: <b>13.0.1601.5</b>	Server Edition: <b>Developer Edition (64-bit)</b>	
<a href="#">Critical/Warning: 0/0</a>	Batch Requests:	Full Scans:
Transactions:	IO Errors: <b>0.0</b>	Packet Errors: <b>1</b>
Page Lookups:	Lock Waits:	Latch Waits:
Worktables Created:	Workfiles Created:	Pending Memory Grants:

---

Last Update: **10-Jun-2019 14:31:31**

## MS SQL Server Performance - HTML

Clicking **Instance Details** in the left/navigation menu opens the **MS SQL Server Performance** display, which allows you to view various database details as well as trending data for the page life expectancy.



## MS SQL Server Waits Table - HTML

Clicking **Wait Stats** in the left/navigation menu opens the **MS SQL Server Waits Table** display, which allows you to view server wait type details in a table format for a particular database server.

MS SQL Server Waits Table 10-Jun-2019 15:07 DATA

Database Connection: win13\_2008express

Wait Types Sorted By Non-Zero Percentage Per Wait Category

Wait Category	Wait Type	Expired	Wait s	Signal s	Resource s
I/O			0.04	0.0	0.04
Lock			0.0	0.0	0.0
Memory			3.97	0.37	3.6
Non-I/O Page Latch			0.82	0.76	0.06
Other			0.0	0.0	0.0
Transaction Log			0.07	0.0	0.07

## MS SQL Server DB Table Sizes - HTML

Clicking **Table Sizes** in the left/navigation menu opens the **MS SQL Server BD Table Sizes** display, which provides database and table size data for a particular SQL database server.

**MS SQL Server DB Table Sizes** 10-Jun-2019 15:30 No Alerts DATA

Database Connection: - All -

**Database Sizes**

Database	Expired	Log Size MB	Row Size MB	State	Total Size MB
model		0.5	1.25	ONLINE	1
tempdb		0.5	2.0	ONLINE	
ReportServer\$SQLEXP		0.75	2.25	ONLINE	
singa		1.0	2.0	ONLINE	

Page 1 of 2 1 - 40 of 46 items

**Table Sizes**

Database	Schema	Table	Row Count	Total Space (MB)	Used Space
jparker	dbo	MSSQL_WAITSTA	0	0.0	
jparker	dbo	MSSQL_SERVERT	0	0.0	
jparker	dbo	MSSQL_SERVERS	0	0.0	
jparker	dbo	MSSQL_PERFCOL	0	0.0	
jparker	dbo	MSSQL_DBSIZES	0	0.0	
jparker	dbo	EMS_TOPICTOTAL	0	0.0	
jparker	dbo	EMS_TOPICSEXT	0	0.0	

## Node.js

The following Node.js Views (and their associated displays) can be found under **Components** tab > **Application/Web Servers** > **Node.js Servers**:

- **"Node/Master View"**: The displays in this View present detailed data for all node instances or for a particular node instance.
- **"Node Request View"**: The displays in this View allow you to view data pertaining to requests for a connection and a host, or view trending request data for a particular URL associated with a connection and a host.
- **"Process View"**: The displays in this View allow you to view the current and historical metrics for all node processes in a heatmap or tabular format for one or all hosts, or view the current and historical metrics for a single node process.

## Node/Master View

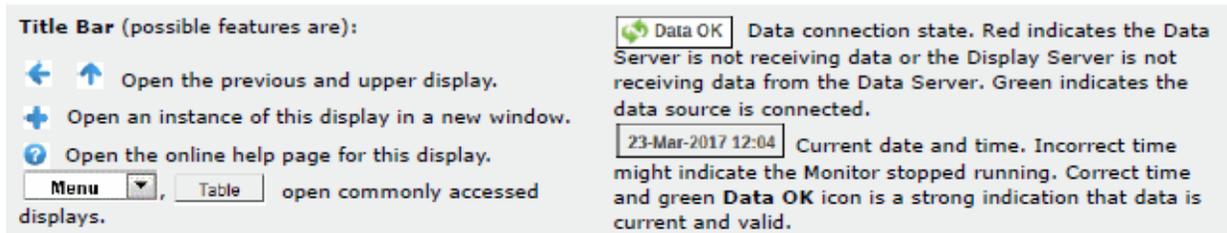
These displays provide detailed data for all node instances or for a particular node instance. Displays in this View are:

- ["Node Master Table"](#): A tabular view of your connected and recently expired node instances and their associated metrics.
- ["Node Master Summary"](#): Provides a way to view trending data for individual node processes.

### Node Master Table

This table provides a view of all your connected (and recently expired) node instances and their associated metric data including host, connection, alert severity, alert count, and the current value of each gathered metric. You can click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by clicking a row to view details for the selected node in the ["Node Master Summary"](#) display.

Connection	Hostname	Alert Level	Alert Count	Process ID	Uptime	CPU %	Request Count	Requests per sec	Requests Delta
SL_Cluster2	TESTBED-29	OK	0	8562	17d 02:50	1.5	11,201,503	0.0	




---

**Note:** The **Requests** button takes you to "[Node Requests Table](#)".

---

### Fields and Data:

<b>Connection</b>	Select the name of the connection containing the node instances for which you want to view data.
<b>Host</b>	Select the name of the host containing the node instances for which you want to view data.
<b>Count</b>	The total number of node instances being monitored based on your search criteria.

### Table:

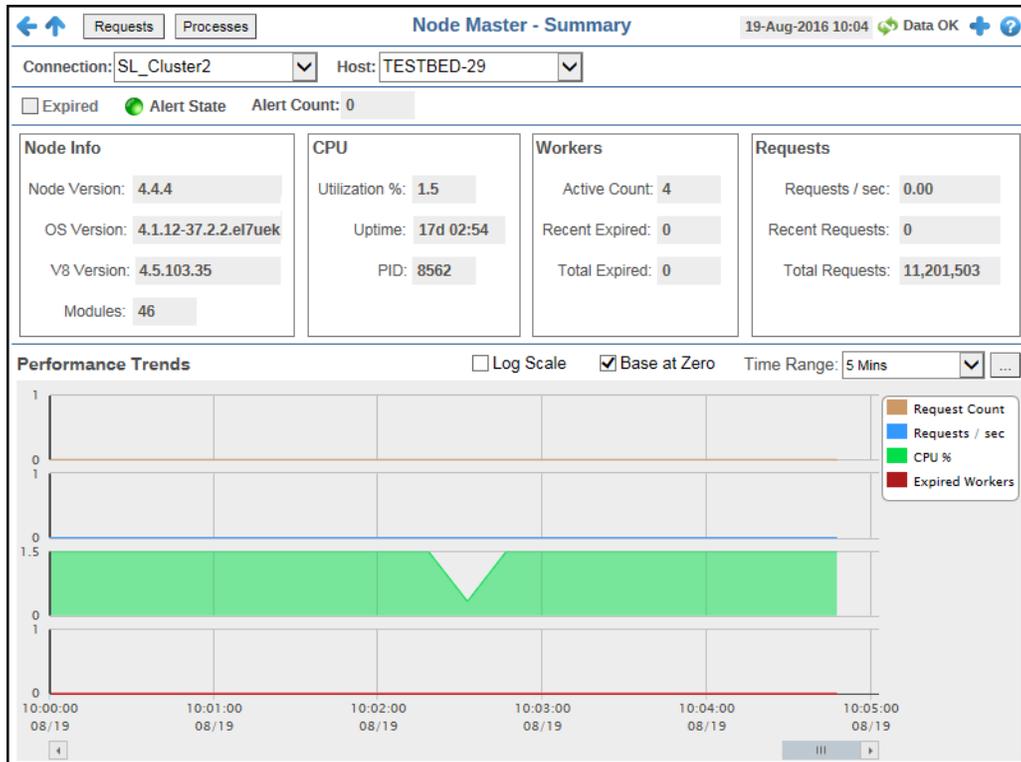
Each row in the table is a different message router.

<b>Connection</b>	The name of the connection.
<b>Host Name</b>	The name of the host.
<b>Alert Level</b>	The current alert severity. <ul style="list-style-type: none"> <li><span style="color: red;">●</span> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li><span style="color: yellow;">●</span> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li><span style="color: green;">●</span> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Alert Count</b>	The total number of current alerts.
<b>Process ID</b>	The process id for the node instance.
<b>Uptime</b>	The amount of time the process has been running.
<b>CPU %</b>	The percentage of CPU used for the process.
<b>Request Count</b>	The total number of requests on the host.
<b>Requests per sec</b>	The average number of requests per second on the host.
<b>Requests Delta</b>	The total number of requests since the last data update.
<b>Requests Mean Rate</b>	The average number of requests for the server since monitoring was started.
<b>Requests 1 Min Rate</b>	The average number of requests for the last minute.
<b>Requests 5 Min Rate</b>	The average number of requests for the last 5 minutes.

<b>Requests 15 min Rate</b>	The average number of requests for the last 15 minutes.
<b>Expired Workers</b>	The number of expired workers on the host since the last data update.
<b>Arch</b>	The CPU architecture of the operating system on the server. Possible values are <b>x64</b> , <b>arm</b> , and <b>ia32</b> .
<b>C-ares</b>	The current version of C-ares running on the host.
<b>Http Parser</b>	The current version of the http parser running on the host.
<b>ICU</b>	The current version of ICU running on the host.
<b>Modules</b>	This number of modules found on the host.
<b>Node Ver</b>	The version of <b>node.js</b> running on the host.
<b>Open SSL</b>	The current version of OpenSSL running on the host.
<b>Platform</b>	The operating system's platform. Possible values, among others, are: <b>darwin</b> , <b>linux</b> , <b>sunos</b> , or <b>win32</b> .
<b>Release</b>	The operating system's release number.
<b>Type</b>	The name of the operating system. Possible values, among others, are <b>Linux</b> on Linux, <b>Darwin</b> on OS X, and <b>Windows_NT</b> on Windows.
<b>UV</b>	The current version of <b>uv</b> running on the host.
<b>V8</b>	The current version of <b>v8</b> running on the host.
<b>ZLib</b>	The current version of <b>ZLib</b> running on the host.
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > ( <b>Project Name</b> ) > <b>Solution Package Configuration</b> > <b>Node.js</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Time Stamp</b>	The date and time the row data was last updated.

## Node Master Summary

This display allows you to view current CPU, worker, and request data as well as trending data for the number of requests, the number of requests per second, the percentage of CPU being used, and the number of recently expired workers on a particular host.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

---

**Note:** The **Requests** button takes you to “[Node Requests Table](#)”. The **Processes** button takes you to “[All \(Node\) Processes Table](#)”.

---

### Filter By:

- Connection** Choose the connection for which you want to show data in the display.
- Host** Choose the host for which you want to show data in the display.

### Fields and Data:

<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > <b>(Project Name)</b> > <b>Solution Package Configuration</b> > <b>Node.js</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Alert State</b>	The current alert severity. <ul style="list-style-type: none"> <li><span style="color: red;">●</span> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li><span style="color: yellow;">●</span> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li><span style="color: green;">●</span> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Alert Count</b>	The total number of current alerts.
<b>Node Info</b>	
<b>Node Version</b>	The version of <b>node.js</b> running on the host.
<b>OS Version</b>	The operating system's version number.
<b>v8 Version</b>	The current version of <b>v8</b> running on the host.
<b>Modules</b>	This number of modules found on the host.
<b>CPU</b>	
<b>Utilization %</b>	The percentage of memory used on the CPU.
<b>Uptime</b>	The amount of time the process has been running.
<b>PID</b>	The process id for the node instance.
<b>Workers</b>	
<b>Active Count</b>	The current number of active workers on the host.
<b>Recent Expired</b>	The number of expired workers on the host since the last data update.
<b>Total Expired</b>	The total number of expired workers on the host.
<b>Requests</b>	
<b>Requests / sec</b>	The average number of requests per second on the host.
<b>Recent Requests</b>	The total number of requests since the last data update.
<b>Total Requests</b>	The total number of requests on the host.
<b>Performance Trends Graph</b>	Traces the following: <ul style="list-style-type: none"> <li><b>Request Count</b> -- traces the number of requests on the host.</li> <li><b>Requests / sec</b> -- traces the number of requests/sec on the host.</li> <li><b>CPU %</b> -- traces the percentage of CPU being used on the host.</li> <li><b>Expired Workers</b> -- traces the number of expired workers on the host.</li> </ul>

- Log Scale** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Base at Zero** Select to use zero (**0**) as the Y axis minimum for all graph traces.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Node Request View

You can view data pertaining to requests for a connection and host, or view trending request data for a particular URL associated with a connection and a host. Displays in this View are:

- **"Node Requests Table"**: A tabular view of request data for one or all hosts on a particular connection.
- **"Node Request Summary"**: Allows you to view trending data (number of requests, number of requests per second, and average response time) for individual URLs by connection and host.

### Node Requests Table

This display allows you to view request data for one or all hosts on a particular connection. You can view the request URL, total number of requests, number of requests per second, the average response time, and the number of recent requests for each host.

Drill-down and investigate by clicking a row to view details for the selected host in the **"Node Request Summary"** display.

Connection	Host	Request URL	Total Requests	Requests Per Sec	A Res Tim
SL_Cluster2	TESTBED-29	/templates/alerts.html	2796103	0.00	
SL_Cluster2	TESTBED-29	/templates/kendo.html	8405400	0.00	

**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

---

**Note:** The **Masters** button takes you to ["Node Master Table"](#).

---

**Filter By:**

The display might include these filtering options:

- Connection**      Select the connection for which you want to view data.
- Host**              Select the host for which you want to view data.

**Fields and Data:**

- Count:**              The total number of nodes (rows) in the table.

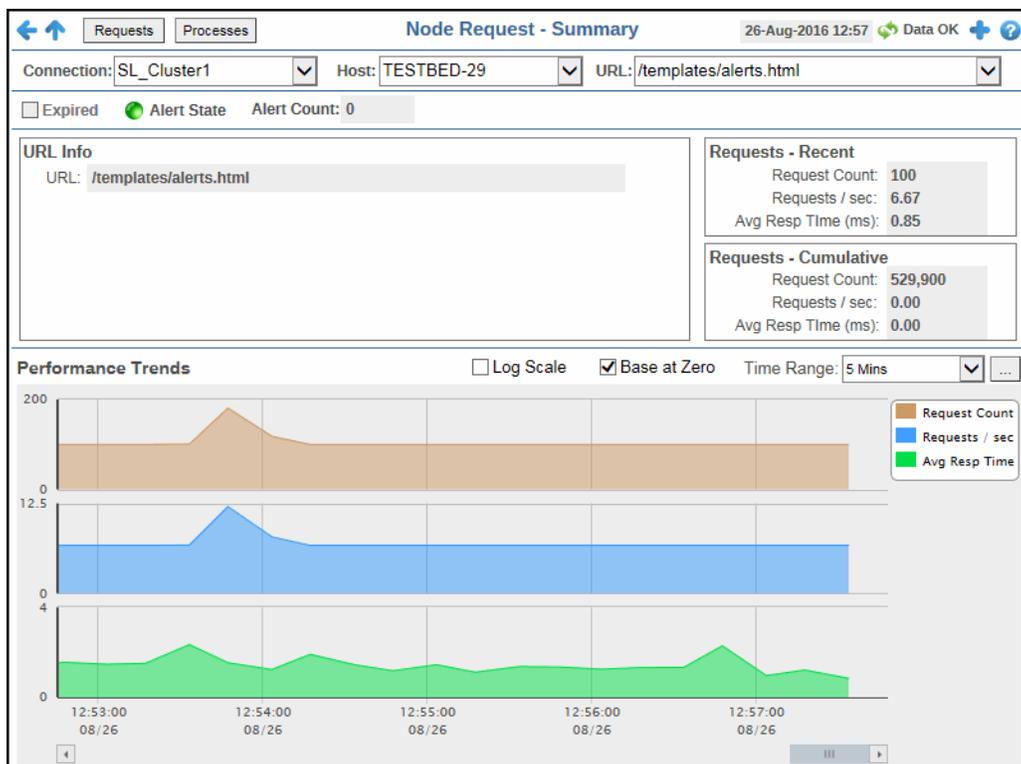
**All Requests Table:**

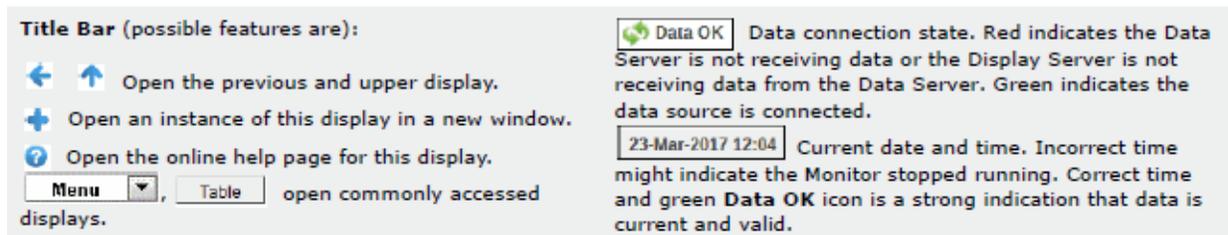
Column values describe the node and its associated requests.

<b>Connection</b>	The name of the connection
<b>Host</b>	The name of the host.
<b>Request URL</b>	The URL from which the requests originated.
<b>Total Requests</b>	The total number of requests.
<b>Requests Per Sec</b>	The rate of requests since the server was started.
<b>Avg Response Time (ms)</b>	The average response time (in milliseconds) since the server was started.
<b>Recent Requests</b>	The total number of requests based on the last query interval.
<b>Recent Requests Per Sec</b>	The rate of recent requests based on the last query interval.
<b>Recent Avg Response Time (ms)</b>	The average response time (in milliseconds) based on the last query interval.
<b>Time Stamp</b>	The date and time the row data was last updated.

## Node Request Summary

This display allows you to view trending data (number of requests, number of requests per second, and average response time) for individual URLs by connection and host.





**Note:** The **Requests** button takes you to “[Node Requests Table](#)”. The **Processes** button takes you to “[All \(Node\) Processes Table](#)”.

### Filter By:

- Connection** Select the connection for which you want to show data in the display.
- Host** Select the host for which you want to show data in the display.
- URL** Select the URL for which you want to view data.

### Fields and Data:

- Expired** When checked, performance data has not been received within the time specified (in seconds) in the **Expire Time** field in the **Duration** region in the RTView Configuration Application > (**Project Name**) > **Solution Package Configuration** > **Node.js** > **DATA STORAGE** tab. The **Delete Time** field (also in the **Duration** region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
- Alert State** The current alert severity.
  - Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
  - Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
  - Green indicates that no metrics have exceeded their alert thresholds.

**Alert Count** The total number of current alerts.

### URL Info

**URL** The URL from which the requests originated.

### Requests - Recent

**Request Count** The total number of requests based on the last query interval.

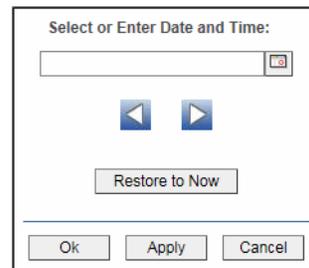
**Requests / sec** The rate of requests based on the last query interval.

**Avg Resp Time (ms)** The average response time (in milliseconds) based on the last query interval.

### Requests - Cumulative

**Request Count** The total number of requests since the server was (re)started.

	<b>Requests / sec</b>	The rate of requests since the server was (re)started.
	<b>Avg Resp Time (ms)</b>	The average response time (in milliseconds) since the server was (re)started.
<b>Performance Trends Graph</b>	Traces the following:	
	<b>Request Count</b>	-- traces the total number of requests.
	<b>Requests / sec</b>	-- traces the rate of requests.
	<b>Avg Resp Time</b>	-- traces the average response time.
	<b>Log Scale</b>	Select to enable a logarithmic scale. Use <b>Log Scale</b> to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. <b>Log Scale</b> makes data on both scales visible by applying logarithmic values rather than actual values to the data.
	<b>Base at Zero</b>	Select to use zero ( <b>0</b> ) as the Y axis minimum for all graph traces.
	<b>Time Range</b>	Select a time range from the drop down menu varying from <b>2 Minutes</b> to <b>Last 7 Days</b> , or display <b>All Data</b> . To specify a time range, click Calendar <input type="button" value="..."/> .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Process View

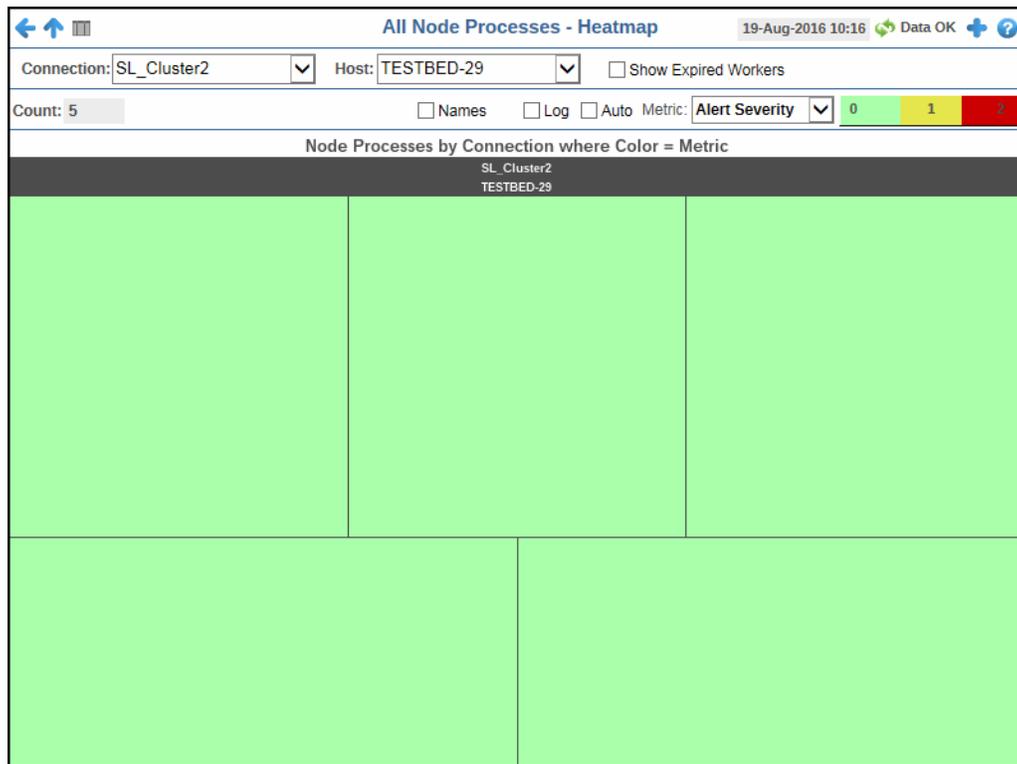
These displays allow you to view the current and historical metrics for all node processes in a heatmap or tabular format for one or all hosts, or view the current and historical metrics for a single node process. Displays in this View are:

- **"All (Node) Processes Heatmap"**: A color-coded heatmap view of data for all node processes for a particular connection/host combination.
- **"All (Node) Processes Table"**: A tabular view of data for all node processes for a particular connection/host combination.
- **"(Node) Process Summary"**: This display allows you to view current and trending data for a single node process for a particular connection/host combination.

## All (Node) Processes Heatmap

This heatmap display provides an easy-to-view interface that allows you to quickly identify the current status of each of your node processes for each available metric. You can view the node processes in the heatmap based on the following metrics: the current alert severity, the current alert count, the percentage of CPU used, and the percentage of memory used. By default, this display shows the heatmap based on the **Alert Severity** metric.

You can use the **Names** check-box  to include or exclude labels in the heatmap, and you can mouse over a rectangle to see additional metrics for a node process. Clicking one of the rectangles in the heatmap opens the “(Node) Process Summary” display, which allows you to see additional details for the selected node process.



### Title Bar (possible features are):

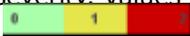
- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

### Fields and Data:

- Connection** Select the connection for which you want to show data in the display.
- Host** Select the host for which you want to show data in the display.

<b>Show Expired Workers</b>	Select this check box to view expired workers in the heatmap.
<b>Count</b>	Lists the total number of processes (rows) found using the search parameters.
<b>Names</b>	Select this check box to display the names of the processes at the top of each rectangle in the heatmap.
<b>Log</b>	Select to enable a logarithmic scale. Use <b>Log Scale</b> to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. <b>Log Scale</b> makes data on both scales visible by applying logarithmic values rather than actual values to the data.
<b>Auto</b>	Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value. <b>Note:</b> Some metrics auto-scale automatically, even when <b>Auto</b> is not selected.
<b>Metric</b>	Choose a metric to view in the display.
<b>Alert Severity</b>	<p>The current alert severity. Values range from <b>0</b> - <b>2</b>, as indicated in the color gradient  bar, where <b>2</b> is the highest Alert Severity:</p> <ul style="list-style-type: none"> <li> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Alert Count</b>	<p>The total number of critical and warning unacknowledged alerts in the instance. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.</p>
<b>CPU Used %</b>	<p>The percentage of CPU used. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of <b>NodeProcessCpuUsageHigh</b>. The middle value in the gradient bar indicates the middle value of the range.</p> <p>When <b>Auto</b> is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>
<b>Memory Used %</b>	<p>The total percentage of memory used. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of <b>NodeProcessMemUsageHigh</b>. The middle value in the gradient bar indicates the middle value of the range.</p> <p>When <b>Auto</b> is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>

## All (Node) Processes Table

This display allows you to view memory, heap memory, and latency data for all processes in a table format. You can drill-down and view the details in the “(Node) Process Summary” display for a specific process by clicking on a row in the resulting table.

Connection	Hostname	Master / Worker	Alert Level	Alert Count	Uptime	CPU %	Process ID	Memory Used (KB)	Memory Used %
SL_Cluster2	TESTBED-29	1	●	0	17d 03:07	0.0	8588	29,905	0.8
SL_Cluster2	TESTBED-29	2	●	0	17d 03:07	0.0	8593	33,124	0.9
SL_Cluster2	TESTBED-29	3	●	0	17d 03:08	0.0	8599	33,042	0.9
SL_Cluster2	TESTBED-29	4	●	0	17d 03:08	0.0	8600	29,815	0.8
SL_Cluster2	TESTBED-29	Master	●	0	17d 03:07	1.5	8562	40,116	1.0

### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

### Filter By:

The display includes these filtering options:

- Connection** Select the connection for which you want to show data in the display.
- Host** Select the host for which you want to show data in the display.
- Show Expired Workers** Select this check box to view expired workers in the table.
- Count** Lists the total number of processes (rows) found using the search parameters.

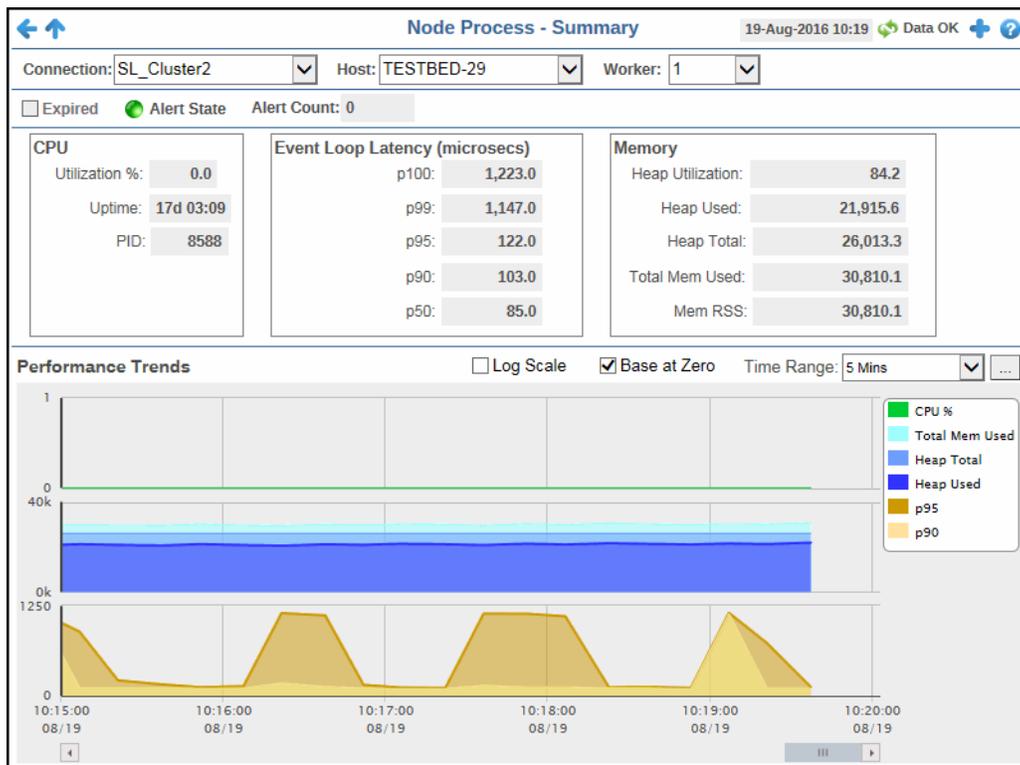
### Fields and Data:

<b>Connection</b>	The name of the connection.
<b>Hostname</b>	The name of the host.
<b>Master / Worker</b>	Displays whether the process is the Master process or, if the application is clustered, the worker ID.
<b>Alert Level</b>	The current alert status.  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
<b>Alert Count</b>	Total number of alerts for the process.
<b>Uptime</b>	Lists the amount of time the process has been up and running.
<b>CPU %</b>	A decimal percentage describing how much the process utilizes the CPU.
<b>Process ID</b>	The process ID.
<b>Memory Used (KB)</b>	The used memory as a fraction of total system memory, in kilobytes.
<b>Memory Used %</b>	The percentage of total available memory used.
<b>Memory RSS (KB)</b>	The Resident Set Size, which is the portion of memory held in RAM (as opposed to swap or disk), in kilobytes.
<b>Heap Total (KB)</b>	The total amount of heap memory from which newly created objects will originate, in kilobytes.
<b>Heap Free (KB)</b>	The amount of memory remaining from which newly created objects will originate, in kilobytes.
<b>Heap Used (KB)</b>	The heap memory currently in use, in kilobytes.
<b>Heap Used %</b>	The percentage of heap memory currently being used.
<b>Heap Avail (KB)</b>	The v8 engine's <b>total_available_size</b> value, in kilobytes.
<b>Heap Limit (KB)</b>	The v8 engine's <b>heap_size_limit</b> value, in kilobytes.
<b>Heap Total Executable (KB)</b>	The v8 engine's <b>total_heap_size_executable</b> value, in kilobytes.
<b>Latency p100</b>	The number of microseconds that 100 percent of events were late in the previous 4 seconds.
<b>Latency p99</b>	The number of microseconds that 99 percent of events were late in the previous 4 seconds.
<b>Latency p95</b>	The number of microseconds that 95 percent of events were late in the previous 4 seconds.
<b>Latency p90</b>	The number of microseconds that 90 percent of events were late in the previous 4 seconds.
<b>Latency p50</b>	The number of microseconds that 50 percent of events were late in the previous 4 seconds.

- Lag** The average number of milliseconds a request has to wait in the Node's event queue before being processed. An excess lag means that the process is overloaded.
- time\_stamp** The date and time the row data was last updated.
- Expired** When checked, performance data has not been received within the time specified (in seconds) in the **Expire Time** field in the **Duration** region in the RTView Configuration Application > (Project Name) > **Solution Package Configuration** > **Node.js** > **DATA STORAGE** tab. The **Delete Time** field (also in the **Duration** region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.

## (Node) Process Summary

This display provides a view of the current and historical metrics for a single process. You can view the current information pertaining to a particular URL and various request data for the node process in the upper portion of the display. The trend graph in the bottom half of the display contains the current and historical number of requests, the number of requests per second, and the average response time for the node process.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Filter By:**

The display might include these filtering options:

<b>Connection</b>	Select the connection for which you want to show data in the display.
<b>Host</b>	Select the host for which you want to show data in the display.
<b>Worker</b>	Select the name of the worker to view. You can select from <b>Master</b> or any of the worker processes created by the Master. Worker processes are defined by numbers: <b>1</b> for the first worker process created by the <b>Master</b> , <b>2</b> for the second worker process created by the <b>Master</b> , and so on.
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > ( <b>Project Name</b> ) > <b>Solution Package Configuration</b> > <b>Node.js</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Alert State</b>	The current alert state of the process.  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
<b>Alert Count</b>	Lists the total number of alerts for the process.
<b>CPU</b>	
<b>Utilization %</b>	A decimal percentage describing how much the process utilizes the CPU.
<b>Uptime</b>	Lists the amount of time the process has been up and running.
<b>PID</b>	The process ID.
<b>Event Loop Latency (microsecs)</b>	
<b>p100</b>	The number of microseconds that 100 percent of events were late in the previous 4 seconds.
<b>p99</b>	The number of microseconds that 99 percent of events were late in the previous 4 seconds.
<b>p95</b>	The number of microseconds that 95 percent of events were late in the previous 4 seconds.
<b>p90</b>	The number of microseconds that 90 percent of events were late in the previous 4 seconds.
<b>p50</b>	The number of microseconds that 50 percent of events were late in the previous 4 seconds.
<b>Memory</b>	
<b>Heap Utilization</b>	The decimal percentage of utilized heap space.
<b>Heap Used</b>	The heap memory currently in use, in kilobytes.
<b>Heap Total</b>	The total amount of memory from which newly created objects can originate, in kilobytes.

<b>Total Mem Used</b>	The used memory as a fraction of total system memory, in kilobytes.
<b>Mem RSS</b>	Resident Set Size, which is the portion of memory held in RAM (as opposed to swap or disk), in kilobytes.

**Performance Trends Graph**

Traces the following:

**CPU %**-- traces the CPU utilization percentage.

**Total Mem Used**-- traces the amount of memory used.

**Heap Total**-- traces the total amount of available heap memory.

**Heap Used**-- traces the amount of used heap memory.

**p95** -- traces the number of microseconds that 95 percent of events were late in the previous 4 seconds.

**p90** -- traces number of microseconds that 90 percent of events were late in the previous 4 seconds.

**Log Scale** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.

**Base at Zero** Select to use zero (**0**) as the Y axis minimum for all graph traces.

**Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

---

## Node.js - HTML

This section describes the HTML version of the Solution Package for Node.js which features an overview display, "[Node.js Overview](#)" (shown below).

The following Node.js Views (and their associated displays) can be found under **Components** tab > **Application/Web Servers**> **Node.js Servers**:

- **"Node.js Overview"**: This display provides a high level overview of the various alerts and metrics for Node.js on a selected data server.
- **"Node/Master View - HTML"**: The displays in this View present detailed data for all node instances or for a particular node instance.
- **"Node Requests View - HTML"**: The displays in this View allow you to view data pertaining to requests for a connection and a host, or view trending request data for a particular URL associated with a connection and a host.
- **"Process View - HTML"**: The displays in this View allow you to view the current and historical metrics for all node processes in a heatmap or tabular format for one or all hosts, or view the current and historical metrics for a single node process.

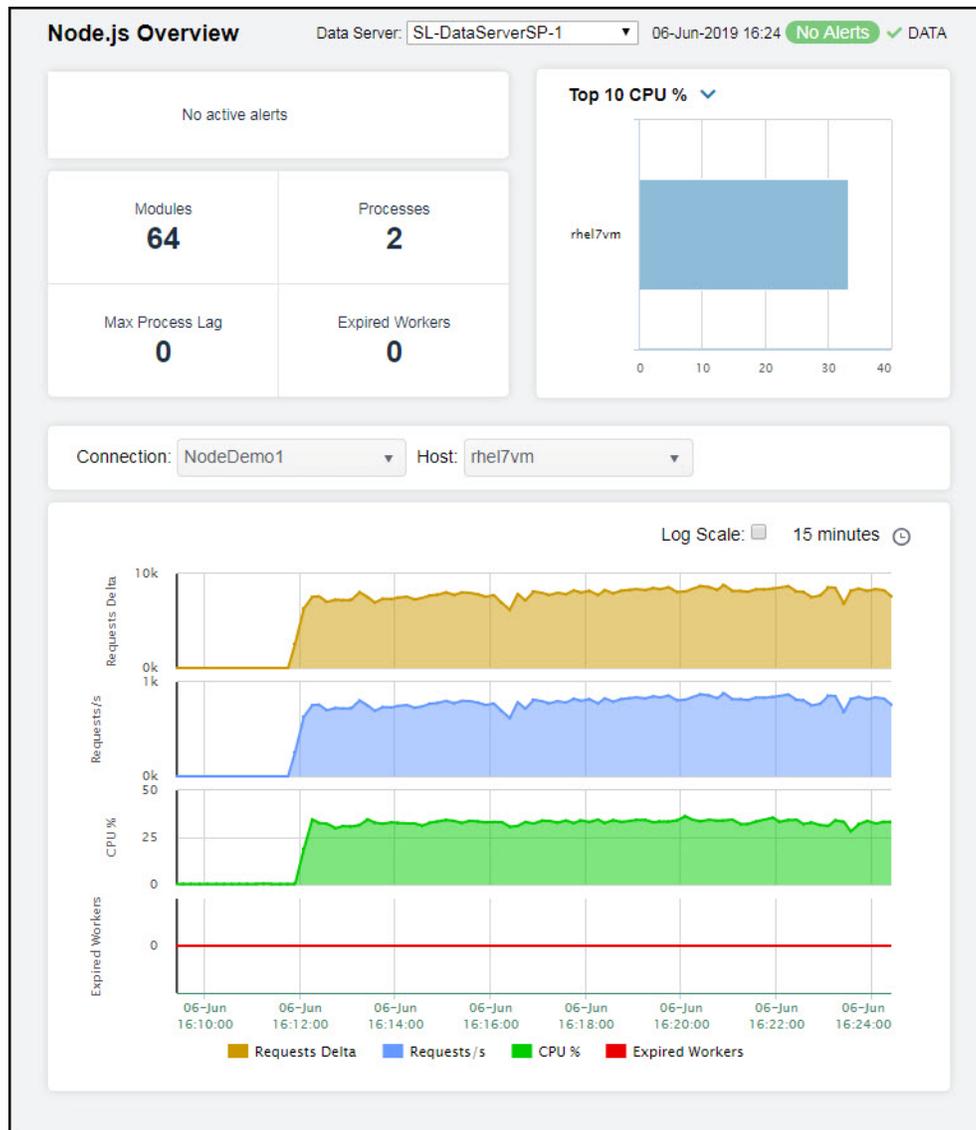
## Node.js Overview

The **Node.js Overview** is the top-level display for the Node.js Monitor, which provides a good starting point for immediately getting the status of all your connections on your Data Server. You can select the RTView DataServer for which you want to see data and easily view the current data for that DataServer including:

- The total number of active alerts for the selected DataServer, including the total number of critical and warning alerts.
- The total number of modules.
- The total number of processes across all instances.
- The maximum process lag across all instances.
- The total number of expired workers.
- A visual list of the top 10 hosts based on CPU percentage and rate of requests on your connected DataServer.

You can hover over each region in the upper half of the Overview to see more detail. You can also drill down to see even more detail by clicking on each respective region in the Overview. For example, clicking on the alerts in the **CRITICAL** and **WARNING** alerts region opens the **Alerts Table by Components** display.

The bottom half of the display provides a trend graph for a selected connection/host combination. You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.



## Node/Master View - HTML

These displays provide detailed data for all node instances or for a particular node instance. Clicking **Node/Master** from the left/navigation menu opens the ["Node/Master Info Table - HTML"](#) display, which shows a tabular view of your connected and recently expired node instances and their associated metrics. The following display is available:

- **Single Node Server:** Opens the ["Node Master Summary - HTML"](#) display, which provides a way to view trending data for individual node processes.

## Node/Master Info Table - HTML

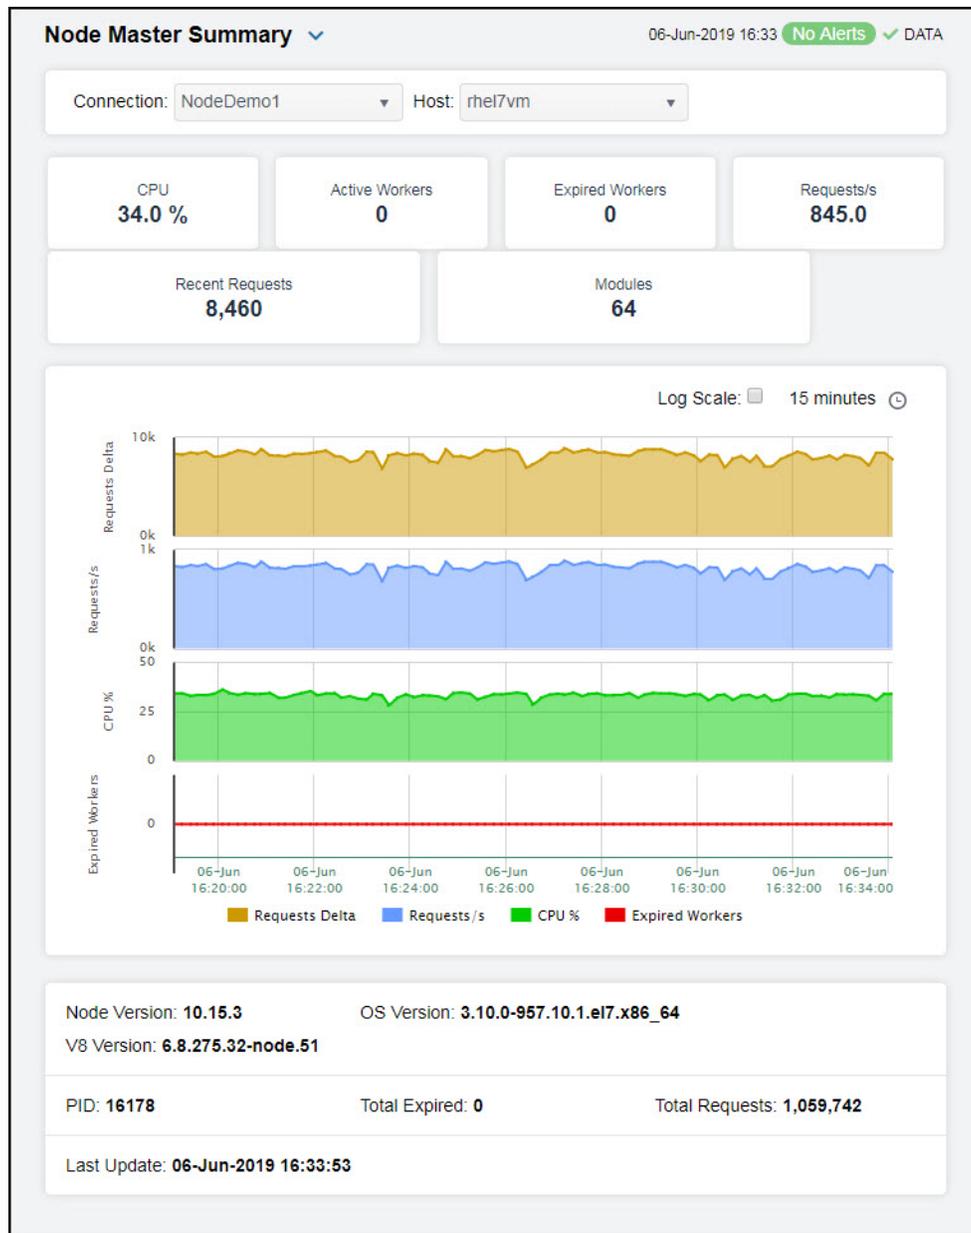
This table provides a view of all your connected (and recently expired) node instances and their associated metric data including host, connection, alert severity, alert count, and the current value of each gathered metric. You can click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by double-clicking a row to view details for the selected node in the ["Node Master Summary - HTML"](#) display. Toggle between the commonly accessed displays by clicking the drop down list on the display title.

Connection	Host	Alert Level	Alert Count	Process ID	Up Ti
NodeDemo1	rhel7vm	✓		16178	40

## Node Master Summary - HTML

Clicking **Single Node Server** in the left/navigation menu opens the **Node Master Summary** display, which allows you to view current CPU, worker, and request data as well as trending data for the number of requests, the number of requests per second, the percentage of CPU being used, and the number of recently expired workers on a particular host. Clicking on the information boxes at the top of the display takes you to the ["Node/Master Info Table - HTML"](#) display, where you can view additional instance data.

Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display.



## Node Requests View - HTML

You can view data pertaining to requests for a connection and host, or view trending request data for a particular URL associated with a connection and a host. Clicking **Node Requests** from the left/navigation menu opens the “[Node Requests Table - HTML](#)” display, which shows a tabular view of request data for one or all hosts on a particular connection. The following display is available:

- **Node Request Summary:** Opens the “[Node Request Summary - HTML](#)” display, which allows you to view trending data (number of requests, number of requests per second, and average response time) for individual URLs by connection and host.

## Node Requests Table - HTML

This display allows you to view request data for one or all hosts on a particular connection. You can view the request URL, total number of requests, number of requests per second, the average response time, and the number of recent requests for each host.

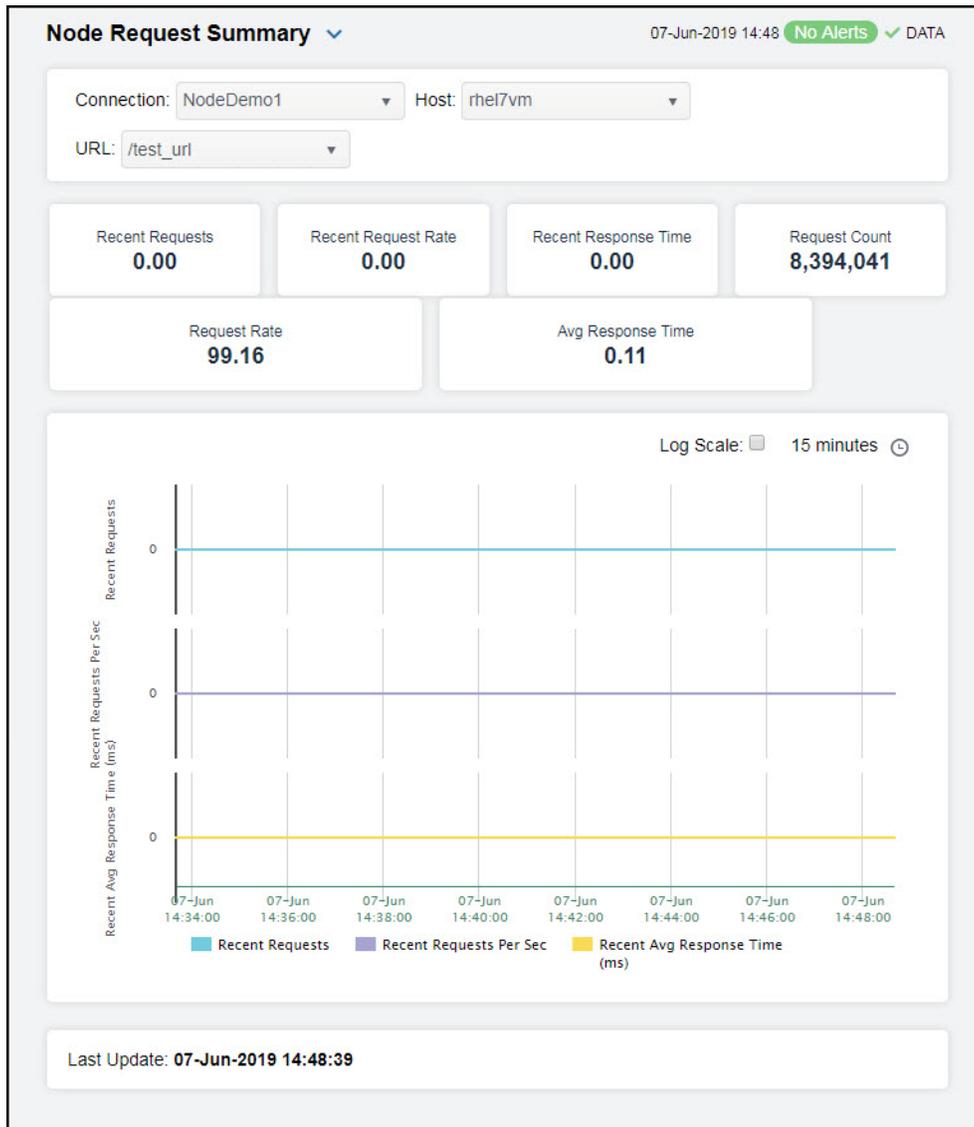
Drill-down and investigate by clicking a row to view details for the selected host in the “[Node Request Summary - HTML](#)” display. Toggle between the commonly accessed displays by clicking the drop down list on the display title.

Connection	Host	Request URL
NODE-HELLOWORLD	192.168.200.201	/

## Node Request Summary - HTML

Clicking **Node Request Summary** in the left/navigation menu opens the **Node Request Summary** display, which allows you to view trending data (number of requests, number of requests per second, and average response time) for individual URLs by connection and host.

Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display.



## Process View - HTML

These displays allow you to view the current and historical metrics for all node processes in a heatmap or tabular format for one or all hosts, or view the current and historical metrics for a single node process. Clicking **Node Processes** from the left/navigation menu opens the ["Node Processes Table - HTML"](#) display, which shows a tabular view of data for all node processes for a particular connection/host combination. The following display is available:

- **Node Processes Heatmap:** Opens the ["Node Processes Heatmap - HTML"](#) display, which consists of a color-coded heatmap view of data for all node processes for a particular connection/host combination.
- **Node Request Summary:** Opens the ["Node Process Summary - HTML"](#) display, which allows you to view current and trending data for a single node process for a particular connection/host combination.

## Node Processes Table - HTML

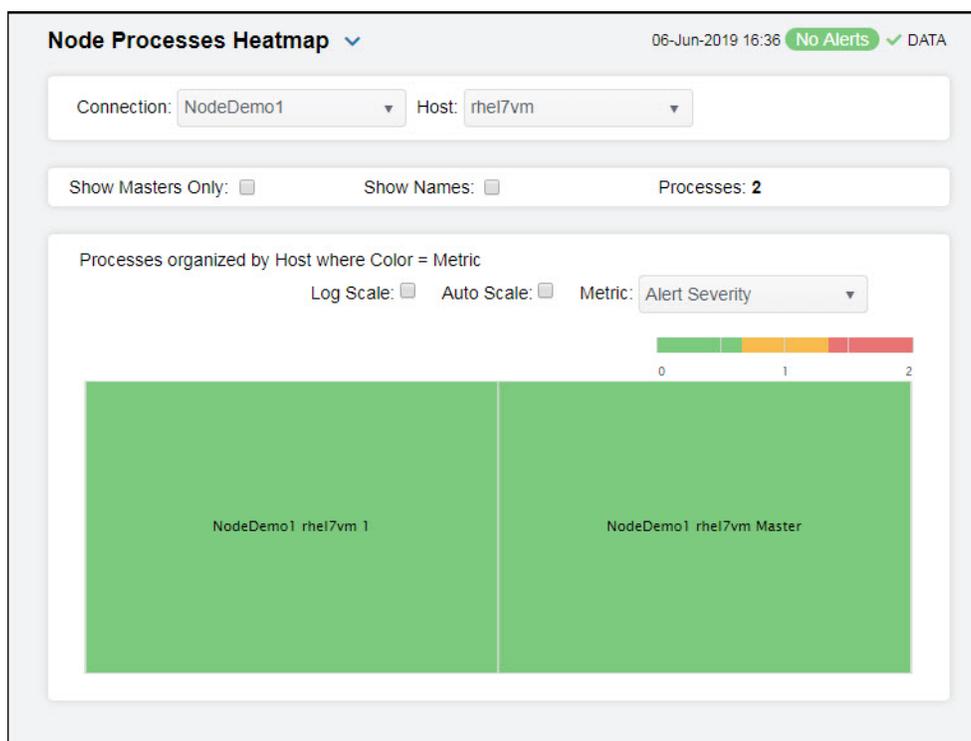
This display allows you to view memory, heap memory, and latency data for all processes in a table format. You can drill-down and view the details in the ["Node Process Summary - HTML"](#) display for a specific process by double-clicking on a row in the resulting table.

Connection	Host	Master/Worker	Alert Level	Alert Count	Up Time
NodeDemo1	rhel7vm	1	✓		1 hr 18 n
NodeDemo1	rhel7vm	Master	✓		1 hr 18 n

## Node Processes Heatmap - HTML

Clicking **Node Processes Heatmap** in the left/navigation menu opens the **Node Processes Heatmap** display, which provides an easy-to-view interface that allows you to quickly identify the current status of each of your node processes for each available metric. You can view the node processes in the heatmap based on the following metrics: the current alert severity, the current alert count, the percentage of CPU used, and the percentage of memory used. By default, this display shows the heatmap based on the **Alert Severity** metric.

You can use the **Show Names** check-box  to include or exclude labels in the heatmap, and you can mouse over a rectangle to see additional metrics for a node process. Clicking **Show Masters Only** results in only the master processes appearing in the heatmap. Clicking one of the rectangles in the heatmap opens the ["Node Process Summary - HTML"](#) display, which allows you to see additional details for the selected node process.



### Available Metrics

Select the metric driving the heatmap display. The default is **Alert Severity**. Each **Metric** has a color gradient bar that maps values to colors. The heatmap organizes the databases by connection, where each rectangle represents a process. Mouse-over any rectangle to display the current values of the metrics for the process. Click on a rectangle to drill-down to the associated ["Node Process Summary - HTML"](#) display for a detailed view of metrics for that particular process.

#### Alert Severity

The current alert severity. Values range from **0 - 2**, as indicated in the color gradient bar, where **2** is the highest Alert Severity:

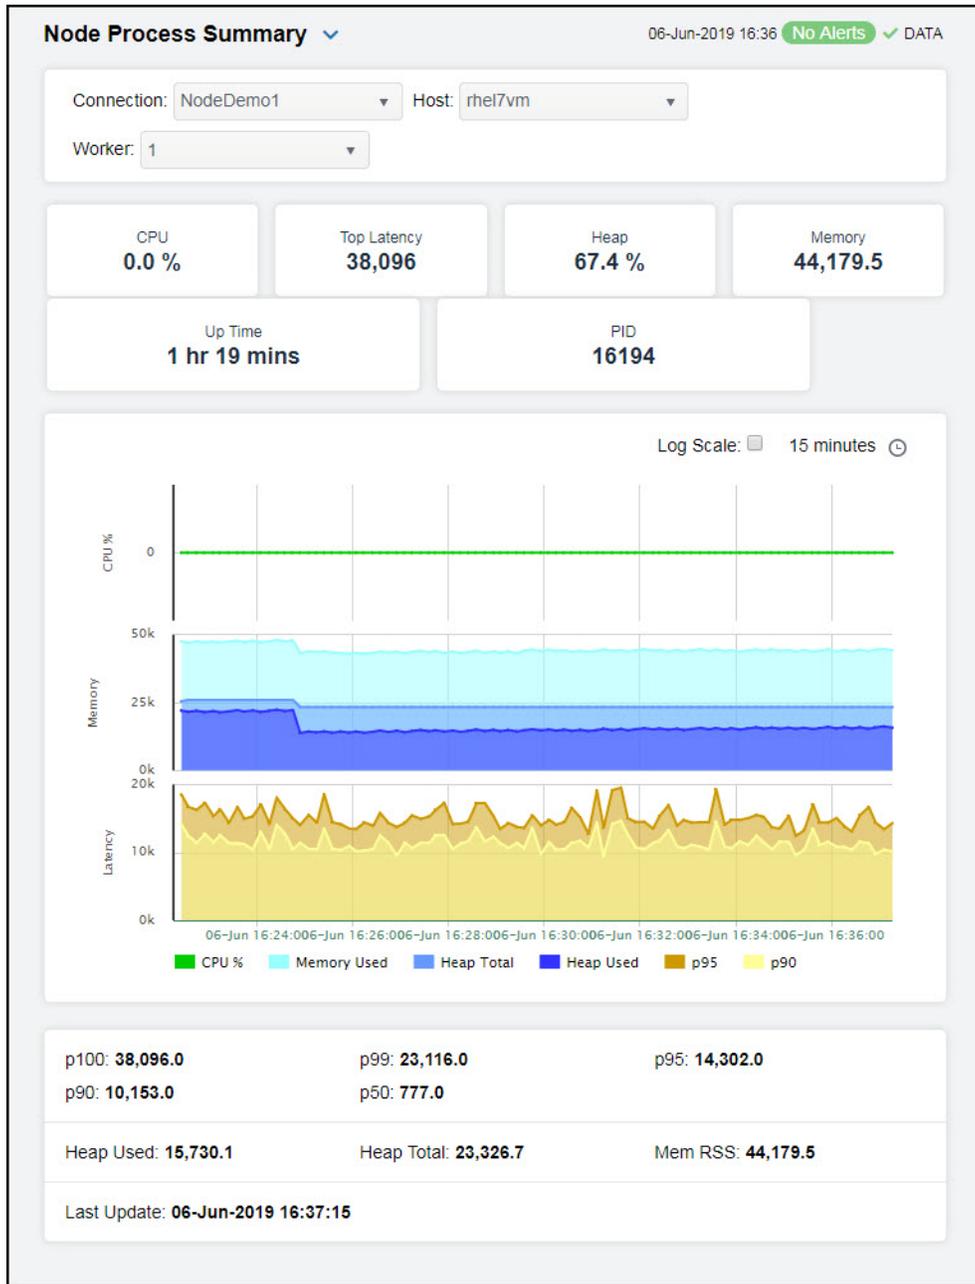
-  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
-  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
-  Green indicates that no metrics have exceeded their alert thresholds.

- Alert Count** The total number of critical and warning unacknowledged alerts in the instance. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.
- CPU Used %** The percentage of CPU used. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **NodeProcessCpuUsageHigh**. The middle value in the gradient bar indicates the middle value of the range.  
When **Auto Scale** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.
- Memory Used %** The total percentage of memory used. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **NodeProcessMemUsageHigh**. The middle value in the gradient bar indicates the middle value of the range.  
When **Auto Scale** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

## Node Process Summary - HTML

Clicking **Node Process Summary** in the left/navigation menu opens the **Node Process Summary** display, which allows you to view the current and historical metrics for a single process. You can view the current information pertaining to a particular process and various request data for the node process in the upper portion of the display. The trend graph in the bottom half of the display contains the current and historical CPU usage percentage, memory usage percentage, heap total, heap used, the number of microseconds that 95 percent of events were late in the previous 4 seconds, and the number of microseconds that 90 percent of events were late in the previous 4 seconds for the node process over a period of time.

Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display.



---

## RTView Host Agent

The Solution Package for RTView Host Agent monitors the health and performance of your physical servers. These predefined displays allow you to be alerted when hosts reach a critical condition. You can also see their performance impact on the technologies and applications they support. Metrics include CPU, memory and storage utilization, process resource consumption and network traffic load. For information about the HTML version, see ["RTView Host Agent - HTML"](#).

The Solution Package for RTView Host Agent is installed onto each host you wish to monitor.

RTView Host Agent displays provide extensive visibility into the health and performance of your hosts. The Solution Package for RTView Host Agent comes with RTView Enterprise. However, the displays are empty until you configure the Solution Package for RTView Host Agent.

RTView Host Agent Views (and their associated displays) can be found under **Components** tab > **Hosts/VMs** > **General Hosts**:

- ["All Hosts"](#)
- ["Single Host"](#)

### All Hosts

These displays present performance data for monitored hosts. Use these displays to examine the state and performance of your hosts. The server displays include summary overviews and detail pages with historical trends.

To see your data in these displays you must install and configure the Solution Package for RTView Host Agent. Displays in this View are:

- ["All Hosts Heatmap"](#)
- ["All Hosts Table"](#)
- ["All Hosts Grid"](#)
- ["All Processes Table"](#)
- ["All Network Table"](#)
- ["All Storage Table"](#)
- ["Host Summary"](#)

### All Hosts Heatmap

View the most critical alert states pertaining to your hosts. Use this display to quickly identify hosts with critical alerts.

Each rectangle in the heatmap represents a host. The rectangle color indicates the most critical alert state associated with the host for the selected **Metric**. The rectangle size represents the amount of physical memory present on the host; a larger size is a larger value.

Choose a domain or **All Domains** from the **Domain** drop-down menu to filter data shown in the display. Choose a different metric to display from the **Metric** drop-down menu. Mouse over a rectangle to see additional metrics. By default, this display shows **Alert Severity**.

Drill-down and investigate a host by clicking a rectangle in the heatmap to view details in the **Host Summary** display.



#### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

#### Filter By:

The display might include these filtering options:

**Domain:** Choose a domain to show data for in the display. Domain names are specified when your administrator configures your Data Server to collect Hawk data, and applies to all host data collected from Hawk by that Data Server.

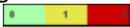
#### Fields and Data:

**Host Count:** The total number of hosts currently shown in the display.

**Show: Domain** When selected, includes the Domain name in the display.

**Host** When selected, includes the Host name in the display.

**Metric** Choose a metric to view in the display.

<b>Alert Severity</b>	<p>The maximum level of alerts in the heatmap rectangle. Values range from <b>0</b> - <b>2</b>, as indicated in the color gradient  bar, where <b>2</b> is the highest Alert Severity:</p> <ul style="list-style-type: none"> <li> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Alert Count</b>	<p>The total number of critical and warning alerts in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.</p>
<b>% CPU Utilization</b>	<p>The percent of CPU used in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.</p>
<b>% Memory Used</b>	<p>The percent of memory used in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.</p>
<b>% Virtual Memory Used</b>	<p>The percent of virtual memory used in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.</p>
<b>1 Minute Load Avg</b>	The average number of processes running over 1 minute.
<b>5 Minute Load Avg</b>	The average number of processes running over 5 minutes.
<b>15 Minute Load Avg</b>	The average number of processes running over 15 minutes.

## All Hosts Table

View host utilization data in a tabular format. Use this display to see all available data for this View.

Each row in the table is a different host. Choose a domain or **All Domains** from the **Domain** drop-down menu. Click a column header to sort column data in numerical or alphabetical order. Drill-down and investigate by clicking a row to view details for the selected application in the **Host Summary** display.

Domain		Host Name	Expired	Severity	Alert Count	Uptime	Host CPU Stats			Memory Used	Memory Total	Memory Used %	Swap Used	Swap Total	Swap Used %	Virtual Us
							% CPU User	% CPU System	% CPU Idle							
myHawkDomain		SLHOST16(sl_amx)	<input type="checkbox"/>	<span style="color: green;">●</span>	0	120d 02:24	8.27	-1.00	91.73	7,309	8,192	89.2	1,581	8,192	19.3	
myHawkDomain		SLHOST16(sl_qa_conn)	<input type="checkbox"/>	<span style="color: green;">●</span>	0	120d 02:21	8.37	-1.00	91.63	7,306	8,192	89.2	1,581	8,192	19.3	
myHawkDomain		SLHOST17(sl_amx)	<input type="checkbox"/>	<span style="color: green;">●</span>	0	120d 02:17	0.71	-1.00	99.29	4,875	8,192	59.5	180	8,192	2.2	
myHawkDomain		SLHOST21(dev)	<input type="checkbox"/>	<span style="color: green;">●</span>	0	120d 04:40	3.03	-1.00	96.97	14,339	16,384	87.5	2,975	16,384	18.2	
myHawkDomain		SLHOST22(sl_qa_conn)	<input type="checkbox"/>	<span style="color: green;">●</span>	0	54d 02:41	0.00	0.00	100.00	2,578	7,824	32.9	0	9,999	0.0	
myHawkDomain		SLHOST5(domain5)	<input type="checkbox"/>	<span style="color: green;">●</span>	0	0d 13:34	17.19	-1.00	82.81	2,313	4,096	56.5	26	4,096	0.6	
myHawkDomain		SLHOST6(domain6)	<input type="checkbox"/>	<span style="color: green;">●</span>	0	0d 13:36	1.87	-1.00	98.13	2,137	3,072	69.6	27	3,072	0.9	

#### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu , Table open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

#### Filter By:

The display might include these filtering options:

**Domain:** Choose a domain to show data for in the display.

#### Fields and Data:

**Host Count:** The total number of hosts in the table.

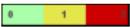
#### Table:

Each row in the table is a different host.

**Domain** The domain in which the host resides. Domain names are specified when your administrator configures your Data Server to collect Hawk data, and applies to all host data collected from Hawk by that Data Server.

**Host Name** The name of the host.

**Expired** When checked, data has not been received from this host in the specified amount of time. The host will be removed from the Monitor in the specified amount of time. The default setting is **60** seconds.

<b>Severity</b>	The maximum level of alerts in the row. Values range from <b>0</b> - <b>2</b> , as indicated in the color gradient  bar, where <b>2</b> is the highest Alert Severity: <ul style="list-style-type: none"> <li><span style="color: red;">●</span> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li><span style="color: yellow;">●</span> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li><span style="color: green;">●</span> Green indicates that no metrics exceeded their alert thresholds.</li> </ul>
<b>Alert Count</b>	The total number of active alerts associated with the host.
<b>Uptime</b>	The amount of time the application has been running, in the following format: <b>0d 00:00 &lt;days&gt;d &lt;hours&gt;:&lt;minutes&gt;:&lt;seconds&gt;</b> For example: <b>10d 08:41:38</b>
<b>% CPU Used</b>	The amount of CPU used, in percent.
<b>% CPU System</b>	The amount of CPU used, in percent.
<b>% CPU Idle</b>	The amount of CPU not used, in percent.
<b>Memory Used</b>	The amount of memory, in megabytes, currently used.
<b>Memory Total</b>	The total amount of memory, in megabytes.
<b>Memory Used%</b>	The amount of memory used, in percent.
<b>Swap Used</b>	The amount of swap space, in megabytes, currently used.
<b>Swap Total</b>	The total amount of swap space, in megabytes.
<b>Swap Used %</b>	The amount of swap space used, in percent.
<b>Virtual Mem(ory) Used</b>	The amount of virtual memory currently used, in megabytes.
<b>Virtual Mem(ory) Total</b>	The total amount of virtual memory, in megabytes.
<b>Virtual Mem(ory) Used%</b>	The amount of virtual memory used, in percent.
<b>Load Avg 1 Minute</b>	The average number of processes running over 1 minute.
<b>Load Avg 5 Minute</b>	The average number of processes running over 5 minutes.
<b>Load Avg 15 Minute</b>	The average number of processes running over 15 minutes.
<b>OS Type</b>	The type of operating system (for example, Linux, HP-UX, Windows 2003).
<b>OS Description</b>	The name of the operating system.
<b>OS Version</b>	The operating system version.
<b>CPU Model</b>	The CPU model.
<b># CPUs</b>	The number of node connections.

<b>Agent Type</b>	The type of agent from which the data was collected: <b>HOSTMON</b> (a SL Host Agent), <b>Hawk</b> , <b>WMI</b> or <b>SNMP</b> .
<b>Agent Class</b>	The specific version of the agent software.
<b>Source</b>	The name of the SL Data Server where the host data was collected.
<b>Timestamp</b>	The date and time the data was last updated.

## All Hosts Grid

This grid provides a list view of utilization metrics for all hosts. Use this display to track and view in parallel the general performance of your hosts. Drill down and investigate by clicking a host to view details in the **Host Summary** display.



### Title Bar (possible features are):

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu, Table open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

### Filter By:

The display might include these filtering options:

<b>Domain:</b>	Choose a domain to show data for in the display. Domain names are specified when your administrator configures your Data Server to collect Hawk data, and applies to all host data collected from Hawk by that Data Server.
<b>Host Count</b>	Displays the number of hosts (including expired hosts) listed in the display.
<b>Time Range:</b>	Choose a time range to show data for in the display. Options are: <b>All Data, 2 Mins, 5 Mins, 20 Mins, 1 Hour, 2 Hours, 4 Hours, 8 Hours, 24 Hours, 2 Days and 7 Days.</b>

**Grid**

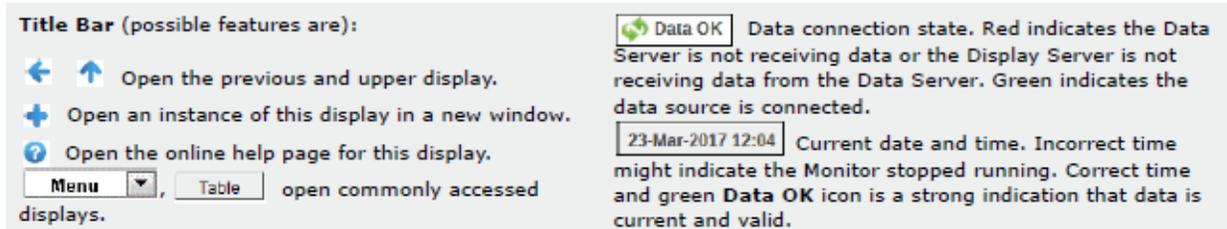
Utilization data shown for hosts in the selected domain.

- Host Name**      The name of the host.
- OS Type**        The name of the operating system.
- Uptime**         The amount of time (days, hours, seconds) the operating system has been running.
- Phys Mem**        The amount of physical memory used, in megabytes.
- Virtual Mem**     The amount of virtual memory used, in megabytes.
- Load Avg**        **1**                The average number of processes running over 1 minute.
- 5**                The average number of processes running over 5 minutes.
- 15**              The average number of processes running over 15 minutes.
- CPU Usage**      The bar graph shows the amount of CPU currently used.
- VMem Usage**     The bar graph shows the amount of virtual memory currently used.
- Trend Graphs**
  - CPU**             Traces the amount of CPU currently used.
  - VM Usage**       Traces the amount of virtual memory currently used.
  - Rx KB/s**         Traces the amount data currently being received per second.
  - Tx KB/s**         Traces the amount data currently being transmitted per second.

**All Processes Table**

View host utilization data in a tabular format. Use this display to see all available data for this View. Each row in the table is a different host. Choose a domain or **All Domains** and a host or **All Hosts** from the drop-down menus. Click a column header to sort column data in numerical or alphabetical order. Drill-down and investigate by clicking a row to view details for the selected application in the **Host Summary** display.

All Processes - Table View											
Domain: All Domains		Host: All Hosts									
Process Count: 687											
Domain	Host Name	Expired	PID	User	Process Name	CPU %	Start Time	Memory Used	Memory Resident	Memory Shared	Page Fault
myHawkDon	SLHOST16(sl_amx)	■	4	<ACCESS DENIE	System	0.02	01-May-2014 23:18:11	17,056	-1	-1	465,4
myHawkDon	SLHOST16(sl_amx)	■	376	NT AUTHORITY\	smss.exe	0.00	01-May-2014 23:18:11	504	-1	-1	1,8
myHawkDon	SLHOST16(sl_amx)	■	540	NT AUTHORITY\	csrss.exe	0.00	01-May-2014 23:18:16	2,472	-1	-1	12,089
myHawkDon	SLHOST16(sl_amx)	■	628	NT AUTHORITY\	wininit.exe	0.00	01-May-2014 23:18:17	172	-1	-1	1,9
myHawkDon	SLHOST16(sl_amx)	■	648	NT AUTHORITY\	csrss.exe	0.00	01-May-2014 23:18:17	216	-1	-1	11,3
myHawkDon	SLHOST16(sl_amx)	■	692	NT AUTHORITY\	services.exe	0.01	01-May-2014 23:18:17	5,736	-1	-1	14,400
myHawkDon	SLHOST16(sl_amx)	■	708	NT AUTHORITY\	lsass.exe	0.02	01-May-2014 23:18:17	9,576	-1	-1	1,273
myHawkDon	SLHOST16(sl_amx)	■	718	NT AUTHORITY\	lsm.exe	0.00	01-May-2014 23:18:17	3,500	-1	-1	1,030
myHawkDon	SLHOST16(sl_amx)	■	800	NT AUTHORITY\	winlogon.exe	0.00	01-May-2014 23:18:17	172	-1	-1	3,6
myHawkDon	SLHOST16(sl_amx)	■	864	<ACCESS DENIE	svchost.exe	0.00	01-May-2014 23:18:20	3,660	-1	-1	1,496
myHawkDon	SLHOST16(sl_amx)	■	416	<ACCESS DENIE	svchost.exe	0.00	01-May-2014 23:18:20	4,376	-1	-1	2,872
myHawkDon	SLHOST16(sl_amx)	■	472	NT AUTHORITY\	LogonUI.exe	0.00	01-May-2014 23:18:21	2,960	-1	-1	164,7
myHawkDon	SLHOST16(sl_amx)	■	640	<ACCESS DENIE	svchost.exe	0.00	01-May-2014 23:18:21	13,756	-1	-1	111,69
myHawkDon	SLHOST16(sl_amx)	■	548	NT AUTHORITY\	svchost.exe	0.05	01-May-2014 23:18:21	121,608	-1	-1	111,21
myHawkDon	SLHOST16(sl_amx)	■	1048	NT AUTHORITY\	svchost.exe	0.28	01-May-2014 23:18:21	26,108	-1	-1	1,605
myHawkDon	SLHOST16(sl_amx)	■	1220	<ACCESS DENIE	svchost.exe	0.00	01-May-2014 23:18:22	7,336	-1	-1	2,716
myHawkDon	SLHOST16(sl_amx)	■	1316	<ACCESS DENIE	svchost.exe	0.00	01-May-2014 23:18:22	13,452	-1	-1	4,123
myHawkDon	SLHOST16(sl_amx)	■	1548	<ACCESS DENIE	spoolsv.exe	0.00	01-May-2014 23:18:23	3,336	-1	-1	434,0
myHawkDon	SLHOST16(sl_amx)	■	1676	<ACCESS DENIE	svchost.exe	0.00	01-May-2014 23:18:23	4,268	-1	-1	3,881
myHawkDon	SLHOST16(sl_amx)	■	1796	NT AUTHORITY\	HeciServer.exe	0.00	01-May-2014 23:18:24	776	-1	-1	12,6
myHawkDon	SLHOST16(sl_amx)	■	1820	NT AUTHORITY\	IProsetMonitor.exe	0.00	01-May-2014 23:18:24	756	-1	-1	10,3
myHawkDon	SLHOST16(sl_amx)	■	2700	<ACCESS DENIE	svchost.exe	0.00	01-May-2014 23:19:05	780	-1	-1	8,8
myHawkDon	SLHOST16(sl_amx)	■	684	<ACCESS DENIE	svchost.exe	0.00	01-May-2014 23:21:06	2,468	-1	-1	2,909
myHawkDon	SLHOST16(sl_amx)	■	2844	NT AUTHORITY\	IASDataMgrSvc.exe	0.00	01-May-2014 23:21:08	5,836	-1	-1	1,102
myHawkDon	SLHOST16(sl_amx)	■	2680	NT AUTHORITY\	jhi_service.exe	0.00	01-May-2014 23:21:19	980	-1	-1	16,6
myHawkDon	SLHOST16(sl_amx)	■	4248	NT AUTHORITY\	SLMS.exe	0.00	01-May-2014 23:21:21	1,724	-1	-1	1,592

**Filter By:**

The display might include these filtering options:

- Domain:** Choose a domain to show data for in the display. Domain names are specified when your administrator configures your Data Server to collect Hawk data, and applies to all host data collected from Hawk by that Data Server.
- Host:** Choose a host to show data for in the display.

**Fields and Data:**

- Process Count:** The total number of processes in the table.

**Table:**

Each row in the table is a different host.

- Domain** The domain in which the host resides.
- Host Name** The name of the host.
- Expired** When checked, data has not been received from this host in the specified amount of time. The host will be removed from the Monitor in the specified amount of time. The default setting is 60 seconds.
- PID** The process ID.
- User** The user name.
- Process Name** The name of the process.
- CPU%** The amount of CPU used, in percent.
- Start Time** The host start time, in the following format:  
**0d 00:00 <days>d <hours>:<minutes>:<seconds>**  
 For example: **10d 08:41:38**
- Memory Used** The amount of memory currently used, in megabytes.
- Memory Resident** The amount of memory currently used by the process that resides in physical memory and is not paged out. Set to **-1** when the data is not available from an agent. (Hawk does not provide this data.)
- Memory Shared** The amount of physical memory that is shared with other processes. Set to **-1** when the data is not available from an agent. (Hawk does not provide this data.)
- Page Faults** The number of page faults.
- Page Faults /sec** The number of page faults per second.
- Timestamp** The date and time the data was last updated.

## All Network Table

View network interface data in a tabular format. Each row in the table is a different network interface card (NIC). Choose a domain or **All Domains** and a host or **All Hosts** from the drop-down menus. Click a column header to sort column data in numerical or alphabetical order.

Interface Count: 4		Host Network Interfaces				
Domain	Host Name	Expired	if Name	Inet Addr	Mask	Flag
QATB	TESTBED-26	<input type="checkbox"/>	lo	127.0.0.1	255.0.0.0	UP LOOPBACK RUNNING
QATB	TESTBED-26	<input type="checkbox"/>	enp0s3	192.168.200.76	255.255.255.0	UP BROADCAST RUNNING
QATB	TESTBED-34	<input type="checkbox"/>	lo	127.0.0.1	255.0.0.0	UP LOOPBACK RUNNING
QATB	TESTBED-34	<input type="checkbox"/>	ens32	192.168.200.34	255.255.255.0	UP BROADCAST RUNNING

### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

### Filter By:

The display might include these filtering options:

- Domain:** Choose a domain for which to show NIC data. Domain names are specified when your administrator configures your Data Server.
- Host:** Choose a host for which to show NIC data.

### Fields and Data:

- Interface Count:** The total number of NICs in the table.

### Table:

Each row in the table is a different NIC.

<b>Domain</b>	The domain in which the NIC resides.
<b>Host Name</b>	The name of the NIC in which the network interface resides.
<b>Expired</b>	When checked, data has not been received from this NIC in the specified amount of time. The NIC will be removed from the Monitor in the specified amount of time. The default setting is 60 seconds.
<b>if Name</b>	The name of the NIC.
<b>Inet Addr</b>	The NIC IP address.
<b>Mask</b>	The NIC subnet mask IP address.
<b>Flags</b>	Descriptive text for NIC flag.
<b>MTU</b>	The largest size packet or frame for the NIC.
<b>Metric</b>	Indicates...
<b>Point To Point</b>	Indicates whether the NIC is a point to point configuration.
<b>Broadcast</b>	Indicates whether the NIC is a broadcast configuration.
<b>rxKBytes</b>	The total number of kilobytes received by the NIC.
<b>rxPackets</b>	The total number of packets received by the NIC.
<b>rxDropped</b>	The total number of received packets that were dropped by the NIC.
<b>rxErrors</b>	The total number of received errors on the NIC.
<b>rxOverruns</b>	The total number of received overruns on the NIC.
<b>rxFrame</b>	The total number of received frames on the NIC.
<b>txKBytes</b>	The total number of kilobytes transmitted by the NIC.
<b>txPackets</b>	The total number of packets transmitted by the NIC.
<b>txDropped</b>	The total number of transmitted packets that were dropped by the NIC.
<b>txErrors</b>	The total number of transmission errors for the NIC.
<b>txOverruns</b>	The total number of transmission overruns for the NIC.
<b>txCollisions</b>	The total number of transmission collisions for the NIC.
<b>txCarrier</b>	The total number of transmission carrier errors for the NIC.
<b>MAC Address</b>	The NIC MAC address.
<b>Rx KB/s</b>	The number of kilobytes received per second.
<b>Tx KB/s</b>	The number of kilobytes transmitted per second.
<b>Rx Packets/s</b>	The number of packets received per second.

**Tx Packets/s** The number of packets transmitted per second.

**Timestamp** The date and time the data was last updated.

## All Storage Table

View storage data in a tabular format. Each row in the table is a different storage partition. Choose a domain or **All Domains** and a host or **All Hosts** from the drop-down menus. Click a column header to sort column data in numerical or alphabetical order.

Domain	Host Name	Expired	File	%	Total	Used	Available	Mount Point	Typ
QATB	WIN-8-CLONE	<input type="checkbox"/>	C:\	86.0	59.90	51.09	8.81	C:\	NTFS/c
QATB	WIN-8-CLONE	<input type="checkbox"/>	\\192.168.200.7	84.0	452.43	377.54	74.89	Z:\	NTFS/r

### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

### Filter By:

The display might include these filtering options:

**Domain:** Choose a domain or **All Domains** to show data for in the display. Domain names are specified when your administrator configures your Data Server to collect Hawk data, and applies to all host data collected from Hawk by that Data Server.

**Host:** Choose a host or **All Hosts** to show data for in the display.

**Fields and Data:**

**Storage Count:** The total number of storage partitions in the table.

**Table:**

Each row in the table is a different host.

<b>Domain</b>	The domain in which the host resides.
<b>Host Name</b>	The name of the host in which the storage partition resides.
<b>Expired</b>	When checked, data has not been received from this host in the specified amount of time. The host will be removed from the Monitor in the specified amount of time. The default setting is 60 seconds.
<b>File System</b>	The storage partition location.
<b>% Used</b>	The amount of storage partition used, in percent.
<b>Total Size (GB)</b>	The storage partition size, in gigabytes.
<b>Used (GB)</b>	The amount of storage partition used, in gigabytes.
<b>Available (GB)</b>	The amount of storage partition available, in gigabytes.
<b>Mount Point</b>	The storage partition parent directory.
<b>Type</b>	The file system type.
<b>Timestamp</b>	The date and time the data was last updated.

## Single Host

These displays present performance data for a single monitored host. Examine details about the health of your hosts.

To see your data in these displays you must install and configure the Solution Package for RTView Host Agent. Displays in this View are:

- ["Host Summary"](#)

## Host Summary

This display provides a detailed view of utilization metrics for a single server.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu Table open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

### Filter By:

The display might include these filtering options:

- Domain:** Choose a domain to show data for in the display. Domain names are specified when your administrator configures your Data Server to collect Hawk data, and applies to all host data collected from Hawk by that Data Server.
- Host:** Choose a host to show data for in the display.
- Expired** When checked, data has not been received from this host in the specified amount of time. The host will be removed from the Monitor in the specified amount of time. The default setting is **60** seconds.
- Last Update** The time the display was last updated.

### Fields and Data:

Data describes the selected host except where noted.

- OS:** The operating system.
- Version:** The operating system version.
- Uptime:** The number of days, hours and minutes since started.
- #CPUs** The number of node connections.

<b>CPU Type:</b>		The type of CPU.
<b>%CPU</b>	<b>User</b>	The amount of CPU used by the user, in percent.
	<b>System</b>	The amount of CPU used by the system, in percent.
	<b>Idle</b>	The amount of CPU that is not used, in percent.
<b>Physical Memory</b>	<b>Used</b>	The amount of physical memory used, in kilobytes.
	<b>Total(MB)</b>	The amount of physical memory available, in kilobytes.
	<b>%Used</b>	The amount of physical memory used, in percent.
<b>Virtual Memory</b>	<b>Used</b>	The amount of virtual memory used, in kilobytes.
	<b>Total(MB)</b>	The amount of virtual memory available, in kilobytes.
	<b>%Used</b>	The amount of virtual memory used, in percent.
<b>Processes</b>		The number of processes running.
<b>Load Avg:</b>	<b>1 Min</b>	The average number of processes running over 1 minute.
	<b>5 Min</b>	The average number of processes running over 5 minutes.
	<b>15 Min</b>	The average number of processes running over 15 minutes.
<b>Storage</b>	<b>File System</b>	The amount of storage space used for the file system, in kilobytes.
	<b>Mount Point</b>	The name used by the operating system to mount and provide an entry point to other storage volumes.
	<b>%Used</b>	The amount of storage space used, in percent.
<b>Network</b>	<b>ifName</b>	The name assigned to the network interface by the operating system.
	<b>RxKB/s</b>	The amount of network data received per second, in kilobytes.
	<b>TxKB/s</b>	The amount of network data transmitted per second, in kilobytes.

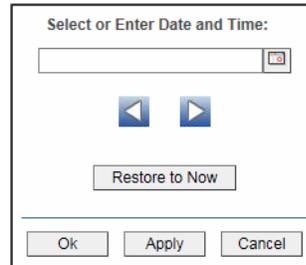
### Trend Graphs

Traces metrics for the selected host.

- **CPU% Used:** The amount of CPU used, in percent.
- **Mem Total:** The amount of available memory, in kilobytes.
- **Mem Used:** The amount of memory used, in kilobytes.
- **Net Rx KB/s:** The amount of network data received per second, in kilobytes.
- **Net Tx KB/s:** The amount of network data transmitted per second, in kilobytes.

**Log Scale** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.

- Base at Zero** Select to use zero (0) as the Y axis minimum for all graph traces.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar [...].



By default, the time range end point is the current time. To change the time range end point, click Calendar [...] and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows ◀ ▶ to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

---

## RTView Host Agent - HTML

The Solution Package for RTView Host Agent monitors the health and performance of your physical servers. These predefined displays allow you to be alerted when hosts reach a critical condition. You can also see their performance impact on the technologies and applications they support. Metrics include CPU, memory and storage utilization, process resource consumption and network traffic load.

the Solution Package for RTView Host Agent is installed onto each host you wish to monitor.

RTView Host Agent displays provide extensive visibility into the health and performance of your hosts. The Solution Package for RTView Host Agent comes with RTView Enterprise. However, the displays are empty until you configure the Solution Package for RTView Host Agent.

The HTML version features an overview display, "[Hosts Overview - HTML](#)", and the following displays which can be found under **Components** tab > **Hosts**:

- "[Hosts Heatmap - HTML](#)"
- "[Hosts Table - HTML](#)"
- "[Host Summary - HTML](#)"
- "[Host Processes - HTML](#)"
- "[Host Network Interfaces - HTML](#)"
- "[Host Storages - HTML](#)"

## Hosts Overview - HTML

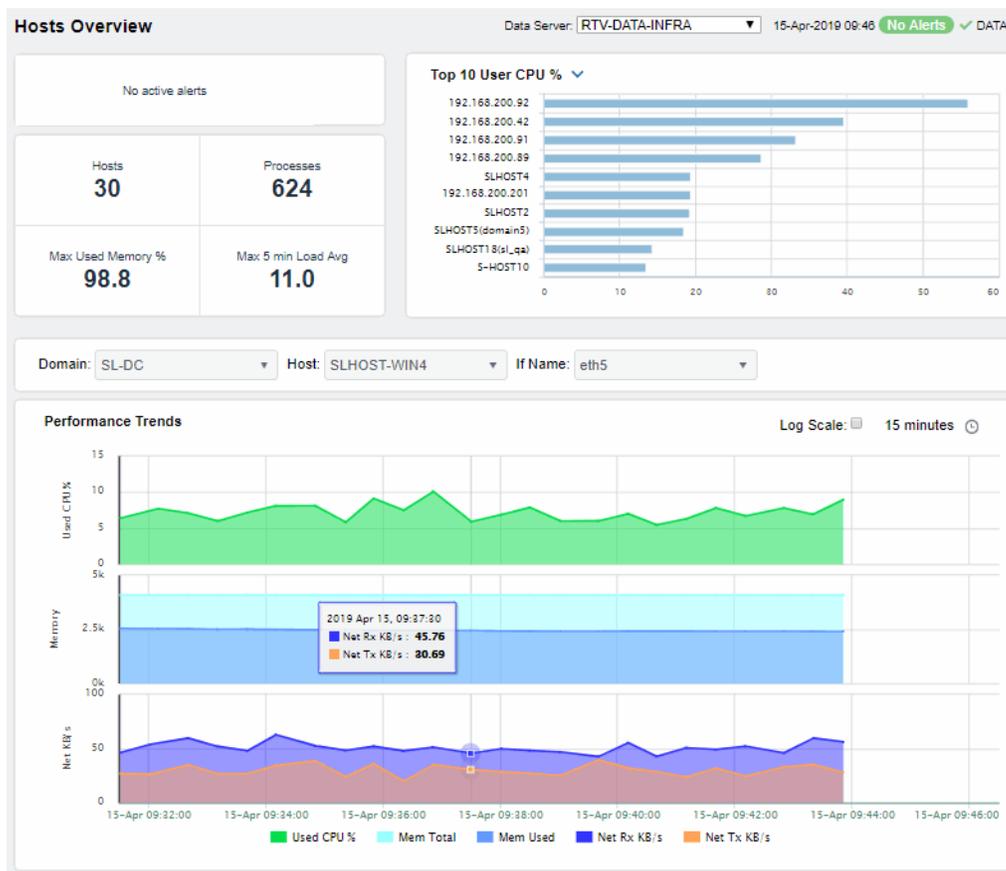
The Hosts Overview is the top-level display for the Hosts Solution Package, which provides a good starting point for immediately getting the status of all your connections on your Data Server. You can select the RTView DataServer for which you want to see data and easily view the current data for that DataServer including:

The total number of active alerts, hosts, processes as well as the maximum percent of memory used and the maximum load average.

A bar graph shows the top 10 CPU users.

You can hover over each region in the upper half of the Overview to see more detail. You can also drill down to see even more detail by clicking on each respective region in the Overview. For example, clicking on the **Max Used Memory %** region opens the ["Host Summary - HTML"](#) display.

The bottom half of the display provides a message rates trend graph for a selected **Domain**, **Host** and **If Name**. You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization in a logarithmic scale and should be used when the range in your data is very broad.



## Hosts Heatmap - HTML

View the status and most critical alert states of all your hosts. Use this display to quickly identify hosts with critical alerts.

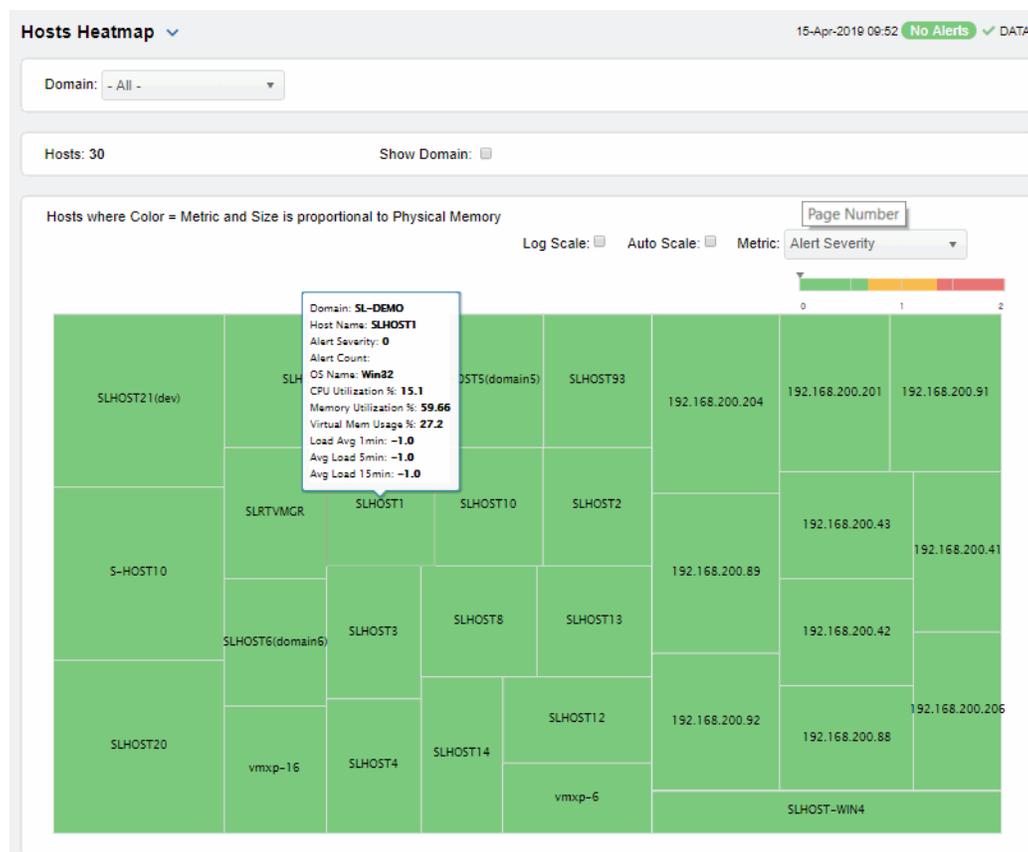
Each rectangle in the heatmap represents a host. The rectangle color indicates the most critical alert state associated with the host for the selected **Metric**. The rectangle size represents the amount of physical memory present on the host; a larger size is a larger value.

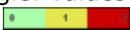
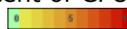
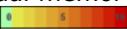
- Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
- Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
- Green indicates that no metrics have exceeded their alert thresholds.

Choose a domain or **All Domains** from the **Domain** drop-down menu to filter data shown in the display. By default, this display shows **Alert Severity**. You can choose other metrics such as **Alert Count**, **CPU Utilization** and **Virtual Memory Utilization**.

Toggle between the commonly accessed Table and Heatmap displays by clicking the drop down list on the display title.

Mouse-over rectangles to view more details about host performance and status. Drill-down and investigate a host by clicking a rectangle in the heatmap to view details in the "[Host Summary - HTML](#)" display.



<b>Metric</b>	Choose a metric to view in the display.
<b>Alert Severity</b>	<p>The maximum level of alerts in the heatmap rectangle. Values range from <b>0</b> - <b>2</b>, as indicated in the color gradient  bar, where <b>2</b> is the highest Alert Severity:</p> <ul style="list-style-type: none"> <li> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Alert Count</b>	<p>The total number of critical and warning alerts in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.</p>
<b>% CPU Utilization</b>	<p>The percent of CPU used in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.</p>
<b>% Memory Used</b>	<p>The percent of memory used in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.</p>
<b>% Virtual Memory Used</b>	<p>The percent of virtual memory used in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.</p>
<b>Load Avg 1 Min</b>	The average number of processes running over 1 minute.

## Hosts Table - HTML

Investigate detailed utilization metrics and configuration details for all hosts in one or all domains. This display contains all metrics available for hosts, including the number of current client connections.

Choose a domain or **All Domains** from the **Domain** drop-down menu. Each row in the table is a different host. Click a column header to sort ascending or descending order. Drill-down and investigate by clicking a row to view details for the selected application in the "[Host Summary - HTML](#)" display. Toggle between the commonly accessed **Table** and **Heatmap** displays by clicking the drop down list on the display title.

Hosts Table 15-Apr-2019 10:17 No Alerts DATA

Domain: - All -

Hosts: 30

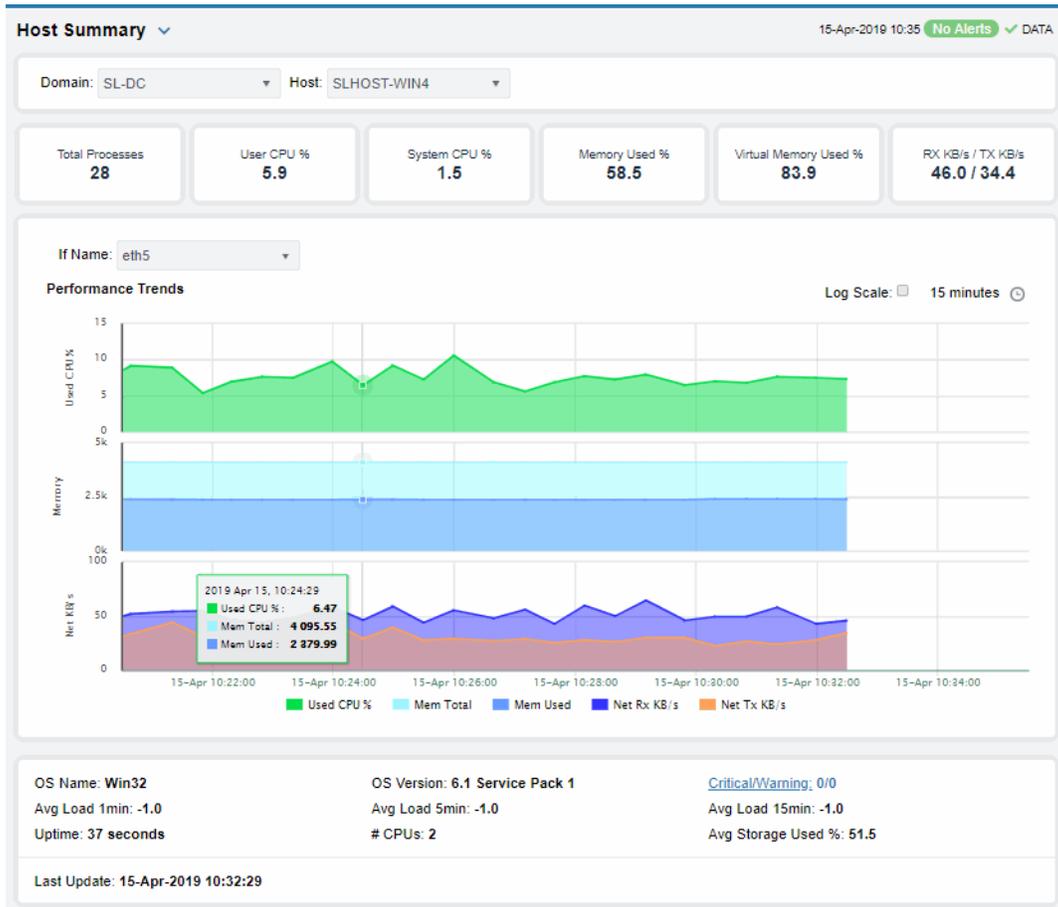
Domain	Host Name	Expired	Alert Level	Alert Count	Uptime	% CPU User	% CPU System	Idle CPU %	Used Memory
SL-DC	SLHOST-WIN4	🕒	✔		36 seconds	7.2	1.5	91.3	2
SL-DEMO	S-HOST10	🕒	✔		36 seconds	14.5	3.2	82.1	5
SL-DEMO	SLHOST1	🕒	✔		15 seconds	10.9	2.9	85.6	1
SL-DEMO	SLHOST10	🕒	✔		37 seconds	2.3	0.6	96.8	2
SL-DEMO	SLHOST12	🕒	✔		36 seconds	10.8	5.1	83.9	1
SL-DEMO	SLHOST13	🕒	✔		36 seconds	1.8	0.2	98.0	1
SL-DEMO	SLHOST14	🕒	✔		15 seconds	11.1	1.2	87.7	2
SL-DEMO	SLHOST18(sl_qa)	🕒	✔		5 hrs 9 mins	18.8	-1.0	81.2	12
SL-DEMO	SLHOST2	🕒	✔		35 seconds	17.4	6.6	74.1	2
SL-DEMO	SLHOST20	🕒	✔		38 seconds	2.5	1.2	96.3	2
SL-DEMO	SLHOST21(dev)	🕒	✔		9 hrs 30 mins	8.5	-1.0	91.5	13
SL-DEMO	SLHOST3	🕒	✔		33 seconds	3.9	3.0	92.8	1
SL-DEMO	SLHOST4	🕒	✔		31 seconds	23.8	7.1	67.7	2
SL-DEMO	SLHOST5(domain5)	🕒	✔		36 seconds	15.1	-1.0	84.9	2
SL-DEMO	SLHOST6(domain6)	🕒	✔		6 hrs 38 mins	1.2	-1.0	98.8	2
SL-DEMO	SLHOST8	🕒	✔		4 hrs 20 mins	0.1	0.2	99.7	2
SL-DEMO	SLHOST93	🕒	✔		3 hrs 6 mins	18.0	1.0	81.0	3
SL-DEMO	SLRTVMGR	🕒	✔		36 seconds	6.2	1.9	91.9	1
SL-DEMO	vmxp-16	🕒	✔		10 hours	0.1	1.2	98.4	1
SL-DEMO	vmxp-6	🕒	✔		9 hrs 18 mins	6.1	3.5	88.2	1
SL-DEMO-LX	192.168.200.201	🕒	✔		14 mins 28 secs	16.8	6.1	71.7	2
SL-DEMO-LX	192.168.200.204	🕒	✔		17 hrs 38 mins	0.7	0.8	98.2	1
SL-DEMO-LX	192.168.200.206	🕒	✔		16 hrs 56 mins	0.5	0.4	98.9	1
SL-DEMO-LX	192.168.200.41	🕒	✔		13 hrs 11 mins	0.8	0.5	97.6	3

## Host Summary - HTML

Track utilization and performance metrics for a single host in any domain. Choose a domain from the **Domain** and a **Host** from the drop-down menus. Clicking on the process and utilization boxes at the top of the display takes you to the "[Hosts Table - HTML](#)" display, where you can view all data available for the host.

In the trend graph region, you can select an **If Name** associated with the selected host and set the time range to trace **Total Processes**, **User CPU%**, **System CPU%**, **Mem Used%**, **Virtual Memory Used%** or **Rx KB/s/TX KB/s**.

Click the **Critical/Warning** link at the bottom of the display to open the Alerts Table by Component display.



## Host Processes - HTML

Investigate detailed process utilization metrics for one or all hosts in one or all domains. This display contains all metrics available for processes including the amount of memory used by each process.

Choose a domain or **All Domains** and a host or **All Hosts** from the drop-down menus. Each row in the table is a different process. Click a column header to sort ascending or descending order.

**Host Processes Table** 15-Apr-2019 10:21 Alerts DATA

Domain: - All - Host: - All -

Processes: 624

Domain	Host Name	Expired	PID	User	Process Name	Proc CPU %	Start Time	Memory Used
SL-DC	SLHOST-WIN4		2448	Administrator	java	0.2	15-Apr-2019 00:05	1,416,110,080
SL-DC	SLHOST-WIN4		3024	Administrator	java:org.hsqldb.Ser	0.02	15-Apr-2019 00:05	857,481,216
SL-DC	SLHOST-WIN4		2880	Administrator	java:com.sl.gmsjau	0.0	15-Apr-2019 00:05	1,355,840,832
SL-DC	SLHOST-WIN4		2744	Administrator	java	0.24	15-Apr-2019 00:05	958,947,328
SL-DC	SLHOST-WIN4		2892	Administrator	java:com.sl.gmsjau	0.0	15-Apr-2019 00:05	1,355,710,464
SL-DC	SLHOST-WIN4		2994	Administrator	java	1.5	15-Apr-2019 00:05	1,524,674,560
SL-DC	SLHOST-WIN4		3096	Administrator	java:com.sl.gmsjau	0.0	15-Apr-2019 00:05	1,356,103,680
SL-DC	SLHOST-WIN4		3788	Administrator	java	0.18	15-Apr-2019 00:05	882,216,192
SL-DC	SLHOST-WIN4		3692	Administrator	java:com.sl.gmsjau	0.0	15-Apr-2019 00:06	1,356,103,680
SL-DC	SLHOST-WIN4		3920	Administrator	java:com.sl.gmsjau	0.0	15-Apr-2019 00:06	1,356,103,680
SL-DC	SLHOST-WIN4		3228	Administrator	java	0.43	15-Apr-2019 00:06	1,513,140,224
SL-DC	SLHOST-WIN4		2376	Administrator	java	0.41	15-Apr-2019 00:06	1,511,043,072
SL-DC	SLHOST-WIN4		3132	Administrator	java:com.sl.gmsjau	0.0	15-Apr-2019 00:06	1,356,103,680
SL-DC	SLHOST-WIN4		4564	Administrator	java	1.31	15-Apr-2019 00:06	1,491,120,128
SL-DC	SLHOST-WIN4		3384	Administrator	java:com.sl.gmsjau	0.0	15-Apr-2019 00:06	1,356,103,680
SL-DC	SLHOST-WIN4		5788	Administrator	java:com.sl.gmsjrtv	0.06	15-Apr-2019 00:06	740,974,592
SL-DC	SLHOST-WIN4		516	Administrator	java:com.sl.gmsjau	0.0	15-Apr-2019 00:06	1,356,103,680
SL-DC	SLHOST-WIN4		604	Administrator	java:com.sl.gmsjrtv	0.27	15-Apr-2019 00:06	1,158,568,964
SL-DC	SLHOST-WIN4		4556	Administrator	java:com.sl.gmsjau	0.0	15-Apr-2019 00:06	1,348,616,192
SL-DC	SLHOST-WIN4		3524	Administrator	java:com.sl.gmsjrtv	1.07	15-Apr-2019 00:06	1,515,216,896
SL-DC	SLHOST-WIN4		5780	Administrator	java:com.sl.gmsjau	0.0	15-Apr-2019 00:06	1,348,616,192
SL-DC	SLHOST-WIN4		5192	Administrator	java:com.sl.gmsjrtv	0.14	15-Apr-2019 00:06	743,071,744

Page 1 of 16 1 - 40 of 624 items

## Host Network Interfaces - HTML

Investigate network interface card (NIC) utilization metrics and configuration details in one or all domains and on one or all hosts. This display contains all metrics available for network interfaces, including IP address, status and RX/TX rates.

Each row in the table is a different NIC. Choose a domain or **All Domains** and a host or **All Hosts** from the drop-down menus. Click a column header to sort ascending or descending order.

**Host Network Interfaces Table** 15-Apr-2019 10:24 Alerts DATA

Domain:  Host:

Interfaces: 66

Domain	Host Name	Expired	If Name	INet Address	Mask	Flags	MTU	Metric
SL-DC	SLHOST-WIN4		eth5	172.16.200.134	255.255.255.0	UP BROADCAST F	1500	0
SL-DC	SLHOST-WIN4		lo0	127.0.0.1	255.0.0.0	UP LOOPBACK RL	1500	0
SL-DEMO	S-HOST10		eth6	192.168.200.167	255.255.255.0	UP BROADCAST F	1500	0
SL-DEMO	S-HOST10		eth9	192.168.220.110	255.255.255.0	UP BROADCAST F	1500	0
SL-DEMO	S-HOST10		lo0	127.0.0.1	255.0.0.0	UP LOOPBACK RL	1500	0
SL-DEMO	SLHOST1		eth0	192.168.200.101	255.255.255.0	UP BROADCAST F	1500	0
SL-DEMO	SLHOST1		lo0	127.0.0.1	255.0.0.0	UP LOOPBACK RL	1520	0
SL-DEMO	SLHOST10		lo0	127.0.0.1	255.0.0.0	UP LOOPBACK RL	1520	0
SL-DEMO	SLHOST10		eth0	192.168.200.110	255.255.255.0	UP BROADCAST F	1500	0
SL-DEMO	SLHOST12		eth0	192.168.200.112	255.255.255.0	UP BROADCAST F	1500	0
SL-DEMO	SLHOST12		lo0	127.0.0.1	255.0.0.0	UP LOOPBACK RL	1520	0
SL-DEMO	SLHOST13		eth6	192.168.200.113	255.255.255.0	UP BROADCAST F	1500	0
SL-DEMO	SLHOST13		lo0	127.0.0.1	255.0.0.0	UP LOOPBACK RL	1500	0
SL-DEMO	SLHOST14		eth6	192.168.200.114	255.255.255.0	UP BROADCAST F	1500	0
SL-DEMO	SLHOST14		lo0	127.0.0.1	255.0.0.0	UP LOOPBACK RL	1500	0
SL-DEMO	SLHOST18(sl_qa)		Intel[R] 82579LM G ?	?	?	?	-1	-1
SL-DEMO	SLHOST18(sl_qa)		Microsoft Loopback ?	?	?	?	-1	-1
SL-DEMO	SLHOST18(sl_qa)		isatap.{427BE387-1 ?	?	?	?	-1	-1
SL-DEMO	SLHOST18(sl_qa)		Teredo Tunneling P ?	?	?	?	-1	-1
SL-DEMO	SLHOST18(sl_qa)		isatap.{3F73C479-? ?	?	?	?	-1	-1
SL-DEMO	SLHOST2		eth0	192.168.200.102	255.255.255.0	UP BROADCAST F	1500	0
SL-DEMO	SLHOST2		lo0	127.0.0.1	255.0.0.0	UP LOOPBACK RL	1520	0
SL-DEMO	SLHOST20		eth6	192.168.200.220	255.255.255.0	UP BROADCAST F	1500	0

Page 1 of 2 1 - 40 of 66 items

## Host Storages - HTML

Investigate storage partition utilization metrics and configuration details in one or all domains and on one or all hosts. This display contains all metrics available for storage partitions including storage size, amount of storage used and status.

Each row in the table is a different storage partition. Choose a domain or **All Domains** and a host or **All Hosts** from the drop-down menus. Click a column header to sort ascending or descending order.

**Host Storages Table** 15-Apr-2019 10:37 Alerts DATA

Domain: SL-DC Host: SLHOST-WIN4

Interfaces: 2

Domain	Host Name	Expired	File System	% Used	Total Size (GB)	Used (GB)	Available (GB)	Mount Point	
SL-DC	SLHOST-WIN4		C:\	52.0	39.9	20.7	19.2	C:\	NTFS
SL-DC	SLHOST-WIN4		\\172.16.200.133\c	51.0	39.9	20.1	19.8	P:\	NTFS

## VMware vCenter

VMware vCenter displays enable you to monitor the health and performance of your virtual machines at the cluster level or the machine level, including viewing datastores, network data events and alerts.

The following Views (and associated displays) can be found under **Components tab > Hosts/VMWare**:

- **"Clusters View"**: View all clusters that are configured on one server or on all servers, and view the high availability and the DRS settings for each of the clusters.
- **"Virtual Machines View"**: View current and historical data for your virtual machines.
- **"Datastores View"**: The displays in this View provide a list of datastores on one or all servers, a list of all hosts mounted to a particular datastore, a list of all virtual machines hosted by a particular datastore, or data for a particular datastore.
- **"Networks View"**: View a list of all networks, as well as data associated with the networks, that exist on one server or on all servers.
- **"Events/Alarms View"**: View event data and alarm data for one server or for all servers.

This section includes the following Views:

- **“Clusters View”**: View all clusters that are configured on one server or on all servers, and view the high availability and the DRS settings for each of the clusters.
- **“Hosts View”**: View performance and utilization data for hosts running on one or all clusters, view utilization data for a specific host running virtual machines, view a list of components contained on a selected host, and view physical and virtual network adapters located on a particular host.
- **“Virtual Machines View”**: View current and historical data for your virtual machines.
- **“Datastores View”**: The displays in this View provide a list of datastores on one or all servers, a list of all hosts mounted to a particular datastore, a list of all virtual machines hosted by a particular datastore, or data for a particular datastore.
- **“Networks View”**: View a list of all networks, as well as data associated with the networks, that exist on one server or on all servers.
- **“Events/Alarms View”**: View event data and alarm data for one server or for all servers.

## Clusters View

The display available in the View lists all clusters that are configured on a particular server or on all servers.

The display available in this view is:

- **“All Clusters”**: View all clusters that are configured on one server or on all servers, and view the high availability and the DRS settings for each of the clusters.

## All Clusters

View all clusters that are configured on one server or on all servers, and view the high availability and the DRS settings for each of the clusters.

← VMware Cluster Compute Resources - Table 22-Mar-2017 10:14 Data OK + ?

Server: All Servers ▾

Cluster Count: 1

		Clusters							
Server	clustname	Alert Severity	Alert Count	Overall Status	# Hosts	# Effective Hosts	HA Enabled		
qavSphere1	SL CORP HA CLUSTER		0		2	2	<input type="checkbox"/>		

**Title Bar** (possible features are):

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu ▾, Table open commonly accessed displays.
- 6,047 The number of items currently in the display.
- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the VMWare vSphere. Refer to VMWare vsphere documentation (<http://pubs.vmware.com/vsphere-65/index.jsp#com.vmware.wssdk.apiref.doc/mo-types-landing.html>) for more information regarding these fields

**Filter By:**

The display might include these filtering options:

- Server** Select the server for which you want to view data.
- Cluster Count** The total number of clusters in the selected server(s), which are listed in the **Clusters** table.

**Clusters Table**

<b>Server</b>	The name of the server.
<b>clustername</b>	The name of the cluster.
<b>Alert Severity</b>	The highest level alert on the cluster.  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
<b>Alert Count</b>	The number of alerts currently on the cluster.
<b>Overall Status</b>	The general health status of the cluster.*  Red indicates that the host is experiencing a problem.  Yellow indicates that the host might have a problem.  Grey indicates that the status of the host's health is unknown.  Green indicates that host's status is OK.
<b># Hosts</b>	Lists the number of hosts on the cluster.*
<b># Effective Hosts</b>	Lists the number of effective hosts.*
<b>HA Enabled</b>	When checked, this check box signifies that High Availability is enabled on the cluster.*
<b>HA Admission Enabled</b>	When checked, this check box signifies that High Availability strict admission is enabled.*
<b>HA Admission Policy</b>	Lists the High Availability admission policy for the cluster.*
<b>HA Datastore Candidate</b>	Displays the High Availability datastore candidate defined on the cluster.*
<b>HA Host Monitoring</b>	Lists whether or not High Availability host monitoring is enabled on the cluster.*
<b>HA VM Monitoring</b>	Lists whether or not High Availability virtual machine monitoring is enabled on the cluster.*
<b>DRS Enabled</b>	When checked, this check box signifies that DRS (Distributed Resource Scheduler) is enabled.*
<b>DRS Enable Behavior Overrides</b>	When checked, this check box signifies that DRS behavior overrides for individual virtual machines are enabled.*
<b>DRS Default VM Behavior</b>	Lists the cluster-wide default DRS behavior for virtual machines.*
<b>Expired</b>	When checked, performance data for that cluster has not been received in the time specified in the <b>Duration</b> region on the RTView Configuration > (Project Name/ <b>MISCMON-LOCAL</b> ) > <b>Solution Package Configuration</b> > <b>VMWare</b> > <b>DATA STORAGE</b> tab.
<b>Timestamp</b>	The date and time the data was last updated.

## Hosts View

These displays present performance and utilization data for hosts running on one or all clusters, present utilization data for a specific host running virtual machines, list components contained on a selected host, and list physical and virtual network adapters for a particular host. Displays in this View are:

- **"All Hosts"**: A tabular view of the utilization data for all hosts running on one or on all clusters.
- **"Single Host Summary"**: Displays utilization data for a specific host running virtual machines.
- **"Host Health"**: View the components contained on a selected host and the component's associated data.
- **"Host NICs"**: View data for all physical and virtual network adapters (NICs) for a particular host.

## All Hosts

View the utilization data for all hosts running on one cluster or on all clusters.

Server	Host Name	Cluster Name	Alert Severity	Alert Count	Overall Status	Connection State	Power State	VMs Hosted
qavSphere1	slesxi-1.sldemos-hq.local	SL CORP HA CLUSTER	<span style="color: green;">●</span>	0	<span style="color: green;">●</span>	connected	poweredOn	26
qavSphere1	slesxi-2.sldemos-hq.local	SL CORP HA CLUSTER	<span style="color: green;">●</span>	0	<span style="color: green;">●</span>	connected	poweredOn	27

**Title Bar** (possible features are):

-   Open the previous and upper display.
-  Open an instance of this display in a new window.
-  Open the online help page for this display.
- open commonly accessed displays.
- The number of items currently in the display.

 **Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

 Open the **Alert Views - RTView Alerts Table** display.

---

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the VMWare vSphere. Refer to VMWare vsphere documentation (<http://pubs.vmware.com/vsphere-65/index.jsp#com.vmware.wssdk.apiref.doc/mo-types-landing.html>) for more information regarding these fields

---

### Filter By:

The display might include these filtering options:

<b>Server</b>	Select the server for which you want to view data.
<b>Cluster</b>	Select the cluster for which you want to view data.
<b>Host Count</b>	The total number of hosts of the selected cluster(s), which are listed in the table

### Hosts Table

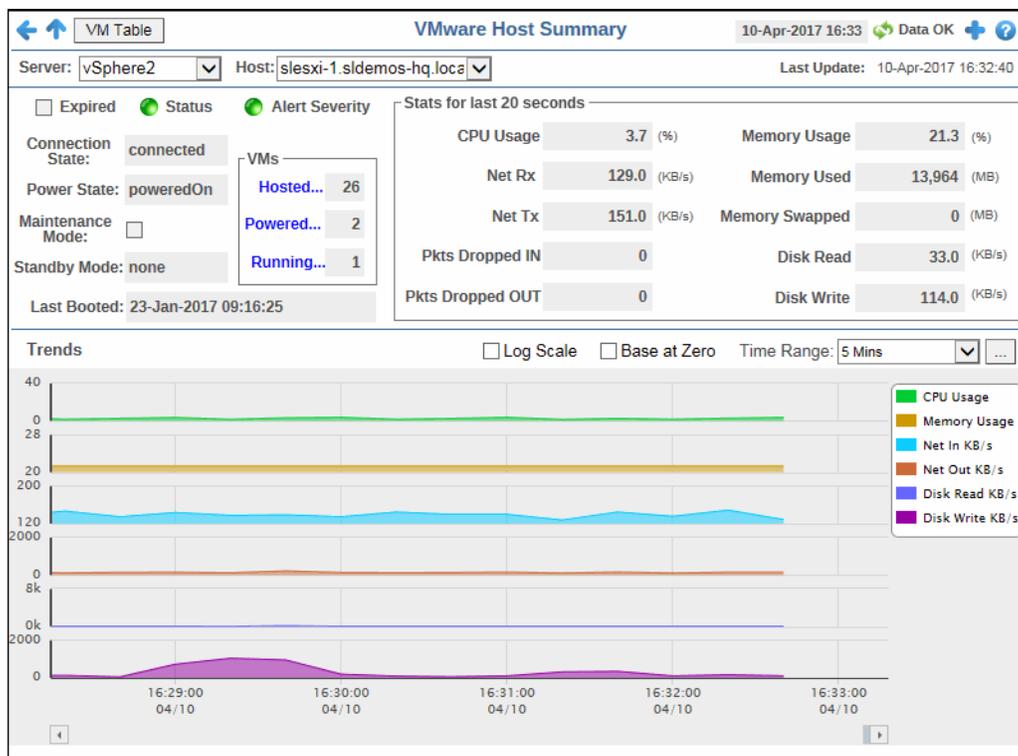
<b>Server</b>	The name of the server.
<b>Host Name</b>	The name of the host.
<b>Cluster Name</b>	The name of the cluster.
<b>Alert Severity</b>	The highest level alert on the host.  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
<b>Alert Count</b>	The number of alerts currently on the host.
<b>Overall Status</b>	The general health status of the host.*  Red indicates that the host is experiencing a problem.  Yellow indicates that the host might have a problem.  Grey indicates that the status of the host's health is unknown.  Green indicates that host's status is OK.
<b>Connection State</b>	Lists the status of the connection.*
<b>Power State</b>	Lists whether the host is powered on or powered off.*
<b>VMs Hosted</b>	The number of virtual machines that exist on the host.*

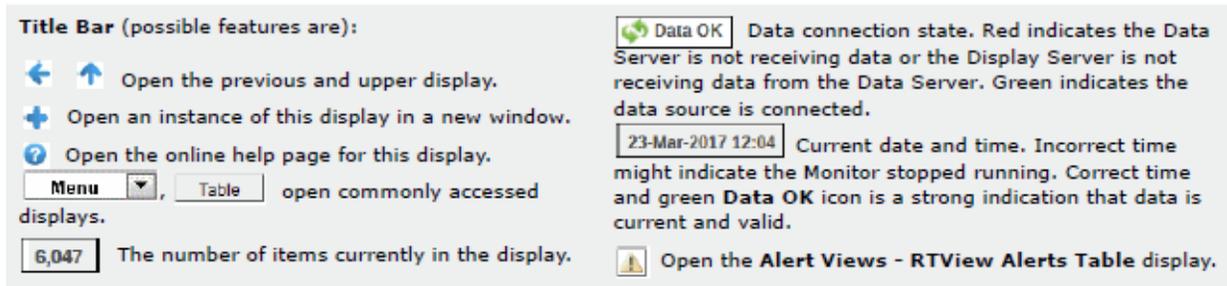
<b>VMs Powered</b>	The number of virtual machines powered on on the host.*
<b>VMs Running</b>	The number of virtual machines running on the host.*
<b>Maintenance Mode</b>	When checked, this check box signifies that the host is in maintenance mode.*
<b>Standby Mode</b>	The host's standby mode.*
<b>CPU % Usage</b>	The percentage of CPU used by the virtual machines.*
<b>Num CPU Cores</b>	The total number of cores on the CPU.*
<b>Num CPU Threads</b>	The total number of threads on the CPU.*
<b>Memory % Usage</b>	The percentage of the host's memory currently in use.*
<b>Memory Used (MB)</b>	The total memory used, in megabytes, on the host.*
<b>Memory Total (MB)</b>	The total amount of memory, in megabytes.*
<b>Swap Used (MB)</b>	The total amount of swap space used by the host, in megabytes.*
<b>Disk Reads (KB/sec)</b>	The amount of data being read from the disk per second, in kilobytes.*
<b>Disk Writes (KB/sec)</b>	The amount of data being written to the disk per second, in kilobytes.*
<b>Net IN (KB/sec)</b>	The amount of network data being received per sec, in kilobytes.*
<b>Net OUT (KB/sec)</b>	The amount of network data being transmitted per sec, in kilobytes.*
<b>% IN Packet Loss (Drops)</b>	The percentage of incoming packets that were dropped.*
<b>% OUT Packet Loss (Drops)</b>	The percentage of outgoing packets that were dropped.*
<b>% IN Packet Loss (Errors)</b>	The percentage of incoming packets that had errors.*
<b>% OUT Packet Loss (Errors)</b>	The percentage of outgoing packets that had errors.*
<b>Packets IN</b>	The number of incoming packets.*
<b>Packets OUT</b>	The number of outgoing packets.*
<b>Packets IN Dropped</b>	The number of incoming packets that were dropped.*
<b>Packets OUT Dropped</b>	The number of outgoing packets that were dropped.*

<b>Packets IN Errors</b>	The number of incoming packets that had errors.*
<b>Packets OUT Errors</b>	The number of outgoing packets that had errors.*
<b>System Vendor</b>	The name of the system vendor.
<b>System Model</b>	The name of the system model.
<b>Expired</b>	When checked, performance data for that cluster has not been received in the time specified in the <b>Duration</b> region on the RTView Configuration > (Project Name/ <b>MISCMON-LOCAL</b> ) > <b>Solution Package Configuration</b> > <b>VMWare</b> > <b>DATA STORAGE</b> tab.
<b>Last Booted</b>	The date and time in which the host was last restarted.*
<b>Timestamp</b>	The date and time the data was last updated.

## Single Host Summary

View the number of virtual machines running on a particular host, the most recent utilization data for the host, and the trend data for the host over a specified time range.





**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the VMWare vSphere. Refer to VMWare vsphere documentation (<http://pubs.vmware.com/vsphere-65/index.jsp#com.vmware.wssdk.apiref.doc/mo-types-landing.html>) for more information regarding these fields.

**Filter By:**

The display might include these filtering options:

- Server** The name of the server containing the host
- Host** The host of the virtual machines for which you want to view data.
- Last Update** The date and time that the data in the table was last updated.

**Fields and Data:**

- Expired** When checked, performance data for that cluster has not been received in the time specified in the **Duration** region on the RTView Configuration > (Project Name/**MISCMON-LOCAL**) > **Solution Package Configuration** > **VMWare** > **DATA STORAGE** tab.
- Status** The general health status of the host.\*
  - Red indicates that the host is experiencing a problem.
  - Yellow indicates that the host might have a problem.
  - Grey indicates that the status of the host’s health is unknown.
  - Green indicates that host’s status is OK.
- Alert Severity** The alert severity for the selected host:
  - Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
  - Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
  - Green indicates that no metrics have exceeded their alert thresholds.
- Connection State** Displays the current state of the connection for the host (**connected/notConnected**).\*
- Power State** Lists whether the host is powered on or powered off.\*
- Maintenance Mode** Signifies whether or not the host is in maintenance mode.\*
  - true:** host is in maintenance mode.
  - false:** host is not in maintenance mode.

<b>Standby Mode</b>	The host's standby mode.*
<b>Last Booted</b>	The date and time in which the host was last restarted.*
<b>VMs</b>	<b>Hosted...</b> The number of virtual machines on the host.*
	<b>Powered...</b> The number of virtual machines on the host that are powered on.*
	<b>Running...</b> The number of virtual machines currently up and running on the host.*
<b>Stats for last 20 seconds</b>	<b>CPU Usage</b> The percentage of CPU used in the last 20 seconds.*
	<b>Net Rx</b> The amount of network data received, in kilobytes per second, in the last 20 seconds.*
	<b>Net Tx</b> The amount of network data transmitted, in kilobytes per second, in the last 20 seconds.*
	<b>Pkts Dropped IN</b> The number of incoming packets that were dropped in the last 20 seconds.*
	<b>Pkts Dropped OUT</b> The number of outgoing packets that were dropped in the last 20 seconds.*
	<b>Memory Usage</b> The percentage of memory used in the last 20 seconds.*
	<b>Memory Used</b> The amount of memory used, in megabytes, in the last 20 seconds.
	<b>Memory Swapped</b> The amount of memory swapped, in megabytes, in the last 20 seconds.*
	<b>Disk Read</b> The amount of data read from the disk, in kilobytes per second, in the last 20 seconds.*
<b>Disk Write</b> The amount of data written to the disk, in kilobytes per second, in the last 20 seconds.*	

### Trend Graphs

Traces the sum of process metrics for the host:

- **CPU Usage:** The percentage of CPU used.
- **Memory Usage:** The amount of memory used.
- **Net In KB/s:** The amount of network data received per second, in kilobytes per second.
- **Net Out KB/s:** The amount of network data transmitted per second, in kilobytes per second.
- **Disk Read KB/s:** The amount of data being read from the disk, in kilobytes per second.
- **Disk Write KB/s:** The amount of data being written to the disk, in kilobytes per second.

**Log Scale** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.

- Base at Zero** Select to use zero (0) as the Y axis minimum for all graph traces.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar  .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**. Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu. Click **Restore to Now** to reset the time range end point to the current time.

## Host Health

View detail data for all the components that are contained on a selected host. Clicking on a component in the table opens details about the associated host in the “Single Host Summary” display.

VMware Host Health - Table				
Server: All Servers		Host: slesxi-1.sldemos-hq.loc		22-Mar-2017 15:42 Data OK
Component Count: 276				
Server	Host Name	Component	Host Health	
Server	Host Name	Component	State	
qavSphere1	slesxi-1.sldemos-hq.local	Broadcom misc-cnic-register 1.72.1.v50.1-1OEM.500.0.0.472560 2012-01-05 01:55:04		
qavSphere1	slesxi-1.sldemos-hq.local	Broadcom net-bnx2 2.2.1j.v50.2-1OEM.500.0.0.472560 2012-04-02 17:31:39.000		
qavSphere1	slesxi-1.sldemos-hq.local	Broadcom net-bnx2x 1.72.18.v50.4-1OEM.500.0.0.472560 2012-04-04 07:36:08.000		
qavSphere1	slesxi-1.sldemos-hq.local	Broadcom net-cnic 1.72.9.v50.1-1OEM.500.0.0.472560 2012-04-02 17:28:46.000		
qavSphere1	slesxi-1.sldemos-hq.local	Broadcom net-tg3 3.123b.v50.1-1OEM.500.0.0.472560 2012-04-03 21:13:10.000		
qavSphere1	slesxi-1.sldemos-hq.local	Broadcom scsi-bnx2fc 1.72.11.v50.1-1OEM.500.0.0.406165 2012-04-02 10:45:54.000		
qavSphere1	slesxi-1.sldemos-hq.local	Broadcom scsi-bnx2i 2.72.10.v50.2-1OEM.500.0.0.472560 2012-04-02 17:30:00.000		
qavSphere1	slesxi-1.sldemos-hq.local	Brocade net-bna 3.0.3.0-1OEM.500.0.0.472560 2011-12-08 08:38:37.000		
qavSphere1	slesxi-1.sldemos-hq.local	Brocade scsi-bfa 3.0.3.0-1OEM.500.0.0.472560 2011-12-08 08:40:51.000		
qavSphere1	slesxi-1.sldemos-hq.local	CPU1 Level-1 Cache is 196608 B		
qavSphere1	slesxi-1.sldemos-hq.local	CPU1 Level-2 Cache is 1572864 B		
qavSphere1	slesxi-1.sldemos-hq.local	CPU1 Level-3 Cache is 15728640 B		
qavSphere1	slesxi-1.sldemos-hq.local	CPU2 Level-1 Cache is 196608 B		
qavSphere1	slesxi-1.sldemos-hq.local	CPU2 Level-2 Cache is 1572864 B		
qavSphere1	slesxi-1.sldemos-hq.local	CPU2 Level-3 Cache is 15728640 B		
qavSphere1	slesxi-1.sldemos-hq.local	Dell dell-configuration-vib 5.0-0 2012-06-18 12:46:14.000		
qavSphere1	slesxi-1.sldemos-hq.local	Dell dell-license-vib 5.0-0 2012-06-18 12:59:33.000		
qavSphere1	slesxi-1.sldemos-hq.local	Dell Inc. BMC Firmware (node 0) 46:10000 1.57		
qavSphere1	slesxi-1.sldemos-hq.local	Dell Inc. System BIOS 2.2.3 2014-05-20 00:00:00.000		
qavSphere1	slesxi-1.sldemos-hq.local	Disk Drive Bay 1 Cable SAS A 0: Config Error - Deassert		
qavSphere1	slesxi-1.sldemos-hq.local	Disk Drive Bay 1 Cable SAS B 0: Config Error - Deassert		
qavSphere1	slesxi-1.sldemos-hq.local	Disk Drive Bay 1 Power Cable 0: Config Error - Deassert		
qavSphere1	slesxi-1.sldemos-hq.local	Disk Drive Bay 1 Signal Cable 0: Config Error - Deassert		
qavSphere1	slesxi-1.sldemos-hq.local	Emulex ima-be2iscsi 4.1.334.3-1OEM.500.0.0.472629 2012-01-11 23:30:20.000		
qavSphere1	slesxi-1.sldemos-hq.local	Emulex net-be2net 4.1.334.0-1OEM.500.0.0.472560 2012-01-09 20:14:01.000		
qavSphere1	slesxi-1.sldemos-hq.local	Emulex scsi-be2iscsi 4.1.334.3-1OEM.500.0.0.472629 2012-01-11 23:30:02.000		
qavSphere1	slesxi-1.sldemos-hq.local	Emulex scsi-lpfc820 8.2.2.126.50-1OEM.500.0.0.472560 2012-01-25 14:35:50.000		
qavSphere1	slesxi-1.sldemos-hq.local	Intel net-lob 3.2.10.1-1OEM.500.0.0.472560 2011-08-22 15:40:14.000		

**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** **Table** open commonly accessed displays.
- 6,047** The number of items currently in the display.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

Open the **Alert Views - RTView Alerts Table** display.

---

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the VMWare vSphere. Refer to VMWare vsphere documentation (<http://pubs.vmware.com/vsphere-65/index.jsp#com.vmware.wssdk.apiref.doc/mo-types-landing.html>) for more information regarding these fields.

---

#### Filter By:

The display might include these filtering options:

<b>Server</b>	The name of the server containing the host
<b>Host</b>	The host of the virtual machines for which you want to view data.
<b>Component Count</b>	The number of components found on the host, which are displayed in the table.

#### Host Health Table

<b>Server</b>	The name of the server.
<b>Host Name</b>	The name of the host.
<b>Component</b>	The name of the component.
<b>State</b>	The general health status of the host.* Red indicates that the host is experiencing a problem. Yellow indicates that the host might have a problem. Grey indicates that the status of the host's health is unknown. Green indicates that host's status is OK.
<b>Sensor Type</b>	Lists the component's sensor type.*
<b>Current Reading</b>	Lists the current reading of the element indicated by the sensor.*
<b>Units</b>	Indicates the base units in which the sensor reading is specified.*
<b>Expired</b>	When checked, performance data for that cluster has not been received in the time specified in the <b>Duration</b> region on the RTView Configuration > (Project Name/ <b>MISCMON-LOCAL</b> ) > <b>Solution Package Configuration</b> > <b>VMWare</b> > <b>DATA STORAGE</b> tab.
<b>Timestamp</b>	The date and time the data was last updated.

## Host NICs

View data for all physical and virtual network adapters (NICs) for a particular host.

VMware Host NICs - Table 22-Mar-2017 15:44 Data OK

Server: All Servers Host: slesxi-1.sldemos-hq.local

Physical NIC Count: 4

Server	Host Name	Device	Link Duplex?	Link Speed (Mb)	AutoNegotiate?	Resource Pool	Scheduler?
qavSphere1	slesxi-1.sldemos-hq.local	vmnic0	<input checked="" type="checkbox"/>	1000	<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>
qavSphere1	slesxi-1.sldemos-hq.local	vmnic1	<input checked="" type="checkbox"/>	1000	<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>
qavSphere1	slesxi-1.sldemos-hq.local	vmnic2	<input checked="" type="checkbox"/>	1000	<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>
qavSphere1	slesxi-1.sldemos-hq.local	vmnic3	<input type="checkbox"/>	0	<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>

Virtual NIC Count: 3

Server	Host Name	Device	DHCP?	IP Address	Subnet Mask	Portgroup
qavSphere1	slesxi-1.sldemos-hq.local	vmk0	<input type="checkbox"/>	192.168.200.51	255.255.255.0	Management Network
qavSphere1	slesxi-1.sldemos-hq.local	vmk1	<input type="checkbox"/>	192.168.201.2	255.255.255.0	vMotion-01
qavSphere1	slesxi-1.sldemos-hq.local	vmk2	<input type="checkbox"/>	192.168.201.3	255.255.255.0	vMotion-02

**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu Table open commonly accessed displays.
- 6,047 The number of items currently in the display.
- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the VMWare vSphere. Refer to VMWare vsphere documentation (<http://pubs.vmware.com/vsphere-65/index.jsp#com.vmware.wssdk.apiref.doc/mo-types-landing.html>) for more information regarding these fields.

**Filter By:**

The display might include these filtering options:

**Server**                      The name of the server containing the host

<b>Host</b>	The host of the virtual machines for which you want to view data.
<b>Physical NIC Count</b>	The number of NICs found on the host, which are displayed in the table.

#### Physical NICs Table

<b>Server</b>	The name of the server.
<b>Host Name</b>	The name of the host.
<b>Device</b>	The name of the device.
<b>Link Duplex?</b>	When checked, indicates that the link is capable of full-duplex. When unchecked, indicates that the link is only capable of half-duplex.*
<b>Link Speed (MB)</b>	The bit rate on the link, in megabytes.*
<b>AutoNegotiate?</b>	When checked, indicates that the physical network adapter supports autonegotiate.*
<b>Resource Pool Scheduler?</b>	When checked, indicates that the physical network adapter allows resource pool-based scheduling for network I/O control.*
<b>VM Direct Path Gen2?</b>	When checked, indicates that the NIC supports VMDirectPath Gen 2.*
<b>Wake On LAN?</b>	When checked, indicates that the NIC is wake-on-LAN capable.*
<b>Driver</b>	The name of the driver.*
<b>MAC</b>	The media access control (MAC) address of the physical network adapter.*
<b>DHCP?</b>	When checked, indicates that the network adapter uses a DHCP server.*
<b>IP Address</b>	The IP address of the physical network adapter.*
<b>Subnet Mask</b>	The subnet mask for the physical network adapter.*
<b>Expired</b>	When checked, performance data for that cluster has not been received in the time specified in the <b>Duration</b> region on the RTView Configuration > (Project Name/ <b>MISCMON-LOCAL</b> ) > <b>Solution Package Configuration</b> > <b>VMWare</b> > <b>DATA STORAGE</b> tab.
<b>Timestamp</b>	The date and time the data was last updated.

#### Virtual NICs Table

<b>Virtual NICs Count</b>	The number of virtual NICs found on the host, which are displayed in the table.
<b>Server</b>	The name of the server.
<b>Host Name</b>	The name of the host.
<b>Device</b>	The name of the device.
<b>DHCP?</b>	When checked, indicates that the network adapter uses a DHCP server.*
<b>IP Address</b>	The IP address of the virtual network adapter.*

<b>Subnet Mask</b>	The subnet mask for the virtual network adapter.*
<b>Port Group</b>	The name of the port group in which the virtual network adapter resides.*
<b>MAC</b>	The media access control (MAC) address of the virtual network adapter.*
<b>TSO Enabled?</b>	When checked, indicates that TCP segment offloading (TSO) is enabled.*
<b>Expired</b>	When checked, performance data for that cluster has not been received in the time specified in the <b>Duration</b> region on the RTView Configuration > (Project Name/ <b>MISCMON-LOCAL</b> ) > <b>Solution Package Configuration</b> > <b>VMWare</b> > <b>DATA STORAGE</b> tab.
<b>Timestamp</b>	The date and time the data was last updated.

## Virtual Machines View

These displays present current and historical data for your virtual machines. Displays in this View are:

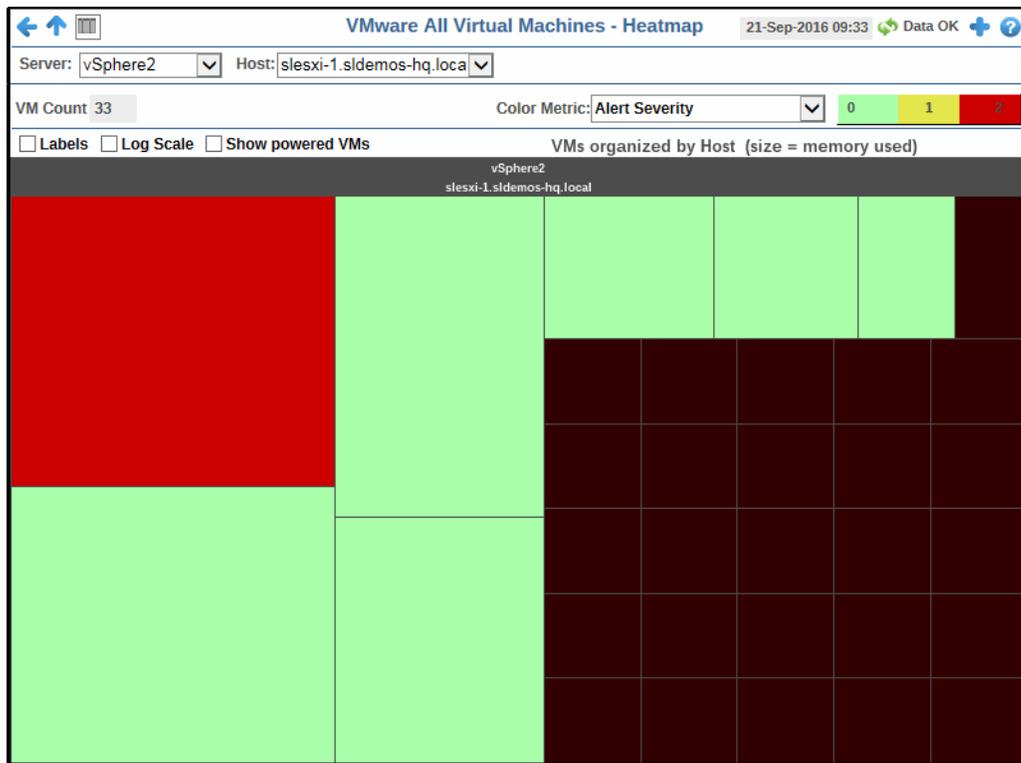
- [“All VMs Heatmap”](#): A color-coded heatmap view of utilization metrics.
- [“All VMs Table”](#): View data shown in the [“All VMs Heatmap”](#) display, as well as additional details, in a tabular format. Use this display to view all available data for each virtual machine by server and host.
- [“All VMs Disk Table”](#): View disk usage percentage, available disk space, and total capacity for one or all virtual machines on a specific server/host combination.
- [“Single VM Summary”](#): View current and historical utilization and performance metrics for a single virtual machine.

### All VMs Heatmap

View the most critical CPU and memory usage, disk read and write utilization, and incoming and outgoing data metrics for your virtual machines. Use this display to quickly identify virtual machines with critical alerts.

Each rectangle in the heatmap represents a virtual machine. The rectangle color indicates the most critical alert state associated with the virtual machine, while the rectangle size represents the maximum memory used in the rectangle (a larger size is a larger value).

Choose a server and host from the drop-down menus to view their associated virtual machines. By default, this display shows **Alert Severity**, but you can choose a different metric to display from the **Color Metric** drop-down menu. Use the **Labels** check-box  to include or exclude labels (virtual machine names for each rectangle) in the heatmap. You can hover your mouse over a rectangle to see additional metrics, and you can drill-down and investigate by clicking a rectangle in the heatmap to view details for the selected application in the [“Single VM Summary”](#) display.



#### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- The number of items currently in the display.

- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

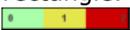
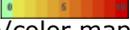
**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the VMWare vSphere. Refer to VMWare vsphere documentation (<http://pubs.vmware.com/vsphere-65/index.jsp#com.vmware.wssdk.apiref.doc/mo-types-landing.html>) for more information regarding these fields

#### Filter By:

The display might include these filtering options:

- Server:** Select the server for which you want to display data.
- Host** Select the host for which you want to display data.
- VM Count:** The total number of virtual machines in the heatmap display.

#### Fields and Data:

<b>Labels</b>	Select this check box to include labels in the heatmap.
<b>Log Scale</b>	Select this check box to enable a logarithmic scale. Use <b>Log Scale</b> to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. <b>Log Scale</b> makes data on both scales visible by applying logarithmic values rather than actual values to the data.
<b>Show powered VMs</b>	Select this check box to include only those VMs that are powered on.
<b>Color Metric</b>	Choose a metric to view in the display.
<b>Alert Severity</b>	<p>The maximum level of alerts in the heatmap rectangle. Values range from <b>0</b> - <b>2</b>, as indicated in the color gradient  bar, where <b>2</b> is the highest Alert Severity:</p> <ul style="list-style-type: none"> <li> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Overall Status</b>	<p>The general health status of the virtual machine.*</p> <ul style="list-style-type: none"> <li> Red indicates that the host is experiencing a problem.</li> <li> Yellow indicates that the host might have a problem.</li> <li> Grey indicates that the status of the host's health is unknown.</li> <li> Green indicates that host's status is OK.</li> </ul>
<b>CPU Usage</b>	The percent (%) CPU used in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Memory Usage</b>	The percent (%) memory used in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Disk Read KB/s</b>	The amount of data being read from the disk per second, in kilobytes, in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Disk Write KB/s</b>	The amount of data being written to the disk per second, in kilobytes, in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Net IN KB/s</b>	The amount of network data received per second, in kilobytes, in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

**Net OUT KB/s** The amount of network data transmitted per second, in kilobytes, in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

**% Net IN Pkts Dropped** The percentage of incoming packets that were dropped in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

**% Net OUT Pkts Dropped** The percentage of outgoing packets that were dropped in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

### All VMs Table

View data shown in the "All VMs Heatmap" display, and additional details, in a tabular format. Use this display to view all available data for each virtual machine by server and host.

Each row in the table lists the details for a virtual machine. Choose a server and a host from the drop-down menus to view all virtual machines running on that server/host combination. You can click a column header to sort column data in numerical or alphabetical order.

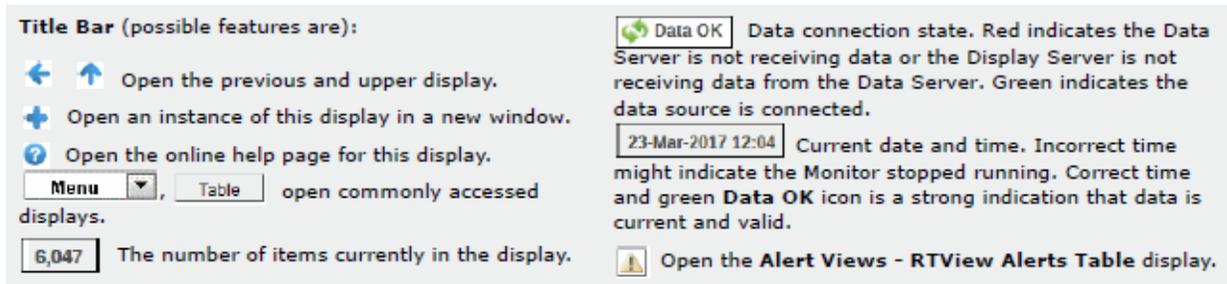
Drill-down and investigate by clicking a row to view details for the selected virtual machine in the "Single VM Summary" display.

VMware Virtual Machines - Table 21-Sep-2016 09:40 Data OK

Server: vSphere2 Host: slesxi-1.sldemos-hq.local

VM Count: 33  Show powered VMs

Server	vSphere VM Name	Alert Severity	Alert Count	Overall Status	Heartbeat	Power State	Guest State	Number CPUs	CPU Usage(%)	Memory Usage(%)
vSphere2	2008-DC		0			poweredOff	notRunning	1	NA	
vSphere2	2008S-IT1		1			poweredOff	notRunning	1	NA	
vSphere2	2008S-WIN10		0			poweredOff	notRunning	2	NA	
vSphere2	2008S-WIN11		0			poweredOff	notRunning	1	NA	
vSphere2	2008S-WIN37		0			poweredOn	running	4	25.0	25.0
vSphere2	2008S-WIN38		0			poweredOff	notRunning	1	NA	
vSphere2	2008S-WIN44		1			poweredOn	running	1	14.7	14.7
vSphere2	64BIT-OL7-0		0			poweredOff	notRunning	2	NA	
vSphere2	64BIT-OL7-1		0			poweredOff	notRunning	2	NA	
vSphere2	64BIT-OL7-10		0			poweredOff	notRunning	2	NA	
vSphere2	64BIT-OL7-15		0			poweredOn	running	2	10.5	10.5
vSphere2	64BIT-OL7-15-BACKUP		0			poweredOff	notRunning	2	NA	
vSphere2	64BIT-OL7-16		0			poweredOn	running	4	3.5	11.0
vSphere2	64BIT-OL7-18		0			poweredOn	running	2	4.8	2.0
vSphere2	64BIT-OL7-2		0			poweredOff	notRunning	2	NA	
vSphere2	64BIT-OL7-3		0			poweredOff	notRunning	2	NA	
vSphere2	64BIT-OL7-5		0			poweredOff	notRunning	2	NA	
vSphere2	64BIT-OL7-6		0			poweredOff	notRunning	2	NA	
vSphere2	64BIT-OL7-8		0			poweredOff	notRunning	2	NA	
vSphere2	64BIT-OL7-9		0			poweredOff	notRunning	2	NA	
vSphere2	64BIT-OL7-UEK4-DOCK...		0			poweredOff	notRunning	4	NA	
vSphere2	64BIT-RH6-3-0		0			poweredOff	notRunning	1	NA	
vSphere2	64BIT-RH6-1		0			poweredOff	notRunning	2	NA	
vSphere2	64BIT-RH6-3-1		0			poweredOn	notRunning	1	8.1	5.0
vSphere2	64BIT-RH6-3-5		0			poweredOff	notRunning	1	NA	
vSphere2	64BIT-RH6-6-9		0			poweredOff	notRunning	2	NA	
vSphere2	64BIT-RHEL7-1		0			poweredOff	notRunning	2	NA	
vSphere2	64BIT-UBUNTU14-1		0			poweredOn	notRunning	4	2.0	
vSphere2	64BIT-W2008S-R2-0		0			poweredOff	notRunning	2	NA	



**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the VMWare vSphere. Refer to VMWare vsphere documentation (<http://pubs.vmware.com/vsphere-65/index.jsp#com.vmware.wssdk.apiref.doc/mo-types-landing.html>) for more information regarding these fields

**Filter By:**

The display might include these filtering options:

- Server:** Select the server containing the virtual machines for which you want to view details.
- Host** Select the host containing the virtual machines for which you want to view details.
- VM Count:** The total number of virtual machines (rows) in the table.
- Show powered VMs** Select to include only those VMs that are powered on.

**Virtual Machines Table:**

Column values describe the virtual machines running on the selected sever/host combination.

- Server** The server on which the virtual machine resides.
- vSphere VM Name** The name of the vSphere virtual machine.
- Alert Severity** The severity of the alert for the virtual machine.
  - Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
  - Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
  - Green indicates that no metrics have exceeded their alert thresholds.
- Alert Count** The total number of active alerts for the virtual machine.
- Overall Status** The general health status of the virtual machine.\*
  - Red indicates that the host is experiencing a problem.
  - Yellow indicates that the host might have a problem.
  - Grey indicates that the status of the host’s health is unknown.
  - Green indicates that host’s status is OK.

<b>Heartbeat</b>	Displays whether or not the virtual machine has a heartbeat.*  Red indicates that heartbeating has stopped.  Grey indicates that heartbeat status is disabled.  Green indicates that heartbeat status is OK.
<b>Power State</b>	Displays whether or not the virtual machine is powered on.*
<b>Guest State</b>	The state of the guest operating system.*
<b>Number CPUs</b>	The number of CPUs used by the virtual machine.*
<b>CPU Usage %</b>	The percentage (%) of CPUs used.*
<b>Memory Usage (%)</b>	The percentage (%) of memory used by the virtual machine.*
<b>Memory Used (MB)</b>	The amount of used memory, in megabytes.*
<b>Memory Total (MB)</b>	The total amount of memory, in megabytes.*
<b>Disk Reads (KB/sec)</b>	The amount of data being read from the disk per second, in kilobytes.*
<b>Disk Writes (KB/sec)</b>	The amount of data being written to the disk per second, in kilobytes.*
<b>Net IN (KB/sec)</b>	The amount of network data received per second, in kilobytes.*
<b>Net OUT (KB/sec)</b>	The amount of network data transmitted per second, in kilobytes.*
<b>% Packet Loss IN</b>	The percentage of incoming packets that have been lost.*
<b>% Packet Loss OUT</b>	The percentage of outgoing packets that have been lost.*
<b>Packets IN</b>	The total number of incoming packets.*
<b>Packets OUT</b>	The total number of outgoing packets.*
<b>Packet IN Dropped</b>	The number of incoming packets that were dropped.*
<b>Packets OUT Dropped</b>	The number of outgoing packets that were dropped.*
<b>Host</b>	The name of the host.*
<b>Guest Host Name</b>	The name of the guest host.*
<b>Guest IP Address</b>	The IP address of the guest.*
<b>Guest Operating System</b>	The operating system used by the guest.*

<b>Connection State</b>	The state of the current connection ( <b>connected/notConnected</b> ).*
<b>Fault Tolerance</b>	Displays whether or not fault tolerance is configured ( <b>configured/notConfigured</b> ).*
<b>VM Tools Run Status</b>	Displays whether or not the guest's tools are running ( <b>guestToolsRunning/guestToolsNotRunning</b> ).*
<b>VM Tools Version Status</b>	Displays the version status of the VMWare tools installed on the guest operating system.*
<b>Expired</b>	When checked, performance data for that cluster has not been received in the time specified in the <b>Duration</b> region on the RTView Configuration > (Project Name/ <b>MISCMON-LOCAL</b> ) > <b>Solution Package Configuration</b> > <b>VMWare</b> > <b>DATA STORAGE</b> tab.
<b>Last Booted</b>	The date and time the virtual machine was last rebooted.*
<b>Timestamp</b>	The date and time the row data was last updated.

### All VMs Disk Table

View disk usage percentage, available disk space, and total capacity for one or all virtual machines on a specific server/host combination.

VMware Virtual Machine Disk Usage - Table 21-Sep-2016 09:45 Data OK

Server: vSphere2 Host: slesxi-1.sldemos-hq.local VM: All VMs

Count: 23

Server	vSphere VM Name	Disk Name	% Disk Usage	Capacity (GB)	Disk Used (GB)	Disk Free (G...	Host
vSphere2	2008-DC	C:\	53.7	20	10.7	9.2	slesxi-1.sldemos-f
vSphere2	2008S-IT1	C:\	95.0	20	18.9	1.0	slesxi-1.sldemos-f
vSphere2	2008S-WIN10	C:\	54.0	30	16.2	13.7	slesxi-1.sldemos-f
vSphere2	2008S-WIN11	C:\	77.1	20	15.3	4.6	slesxi-1.sldemos-f
vSphere2	2008S-WIN37	C:\	77.0	30	23.0	6.9	slesxi-1.sldemos-f
vSphere2	2008S-WIN44	C:\	95.1	50	47.4	2.5	slesxi-1.sldemos-f
vSphere2	64BIT-OL7-16	/	52.3	17	9.2	8.3	slesxi-1.sldemos-f
vSphere2	64BIT-OL7-16	/boot	57.9	0	0.3	0.2	slesxi-1.sldemos-f
vSphere2	64BIT-OL7-16	/tmp	52.3	17	9.2	8.3	slesxi-1.sldemos-f
vSphere2	64BIT-OL7-16	/var/tmp	52.3	17	9.2	8.3	slesxi-1.sldemos-f
vSphere2	64BIT-OL7-18	/	68.6	17	12.0	5.5	slesxi-1.sldemos-f
vSphere2	64BIT-OL7-18	/boot	57.9	0	0.3	0.2	slesxi-1.sldemos-f
vSphere2	64BIT-OL7-18	/tmp	68.6	17	12.0	5.5	slesxi-1.sldemos-f
vSphere2	64BIT-OL7-18	/var/tmp	68.6	17	12.0	5.5	slesxi-1.sldemos-f
vSphere2	64BIT-OL7-6	/	40.6	17	7.1	10.4	slesxi-1.sldemos-f
vSphere2	64BIT-OL7-6	/boot	57.4	0	0.3	0.2	slesxi-1.sldemos-f
vSphere2	64BIT-OL7-6	/tmp	40.6	17	7.1	10.4	slesxi-1.sldemos-f
vSphere2	64BIT-OL7-6	/var/tmp	40.6	17	7.1	10.4	slesxi-1.sldemos-f
vSphere2	64BIT-OL7-UEK4-DOCKER1....	/	41.2	17	7.2	10.3	slesxi-1.sldemos-f
vSphere2	64BIT-OL7-UEK4-DOCKER1....	/boot	57.9	0	0.3	0.2	slesxi-1.sldemos-f
vSphere2	64BIT-OL7-UEK4-DOCKER1....	/tmp	41.2	17	7.2	10.3	slesxi-1.sldemos-f
vSphere2	64BIT-OL7-UEK4-DOCKER1....	/var/tmp	41.2	17	7.2	10.3	slesxi-1.sldemos-f
vSphere2	QAWIN3	C:\	80.3	25	20.0	4.9	slesxi-1.sldemos-f

**Title Bar** (possible features are):

-   Open the previous and upper display.
-  Open an instance of this display in a new window.
-  Open the online help page for this display.
- open commonly accessed displays.
- The number of items currently in the display.

 **Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

 Open the **Alert Views - RTView Alerts Table** display.

---

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the VMWare vSphere. Refer to VMWare vsphere documentation (<http://pubs.vmware.com/vsphere-65/index.jsp#com.vmware.wssdk.apiref.doc/mo-types-landing.html>) for more information regarding these fields

---

#### Filter By:

The display might include these filtering options:

<b>Server</b>	Select the server containing the virtual machine(s) for which you want to view data.
<b>Host</b>	Select the host containing the virtual machine(s) for which you want to view data.
<b>VM</b>	Select a virtual machine for which you want to view data, or select <b>All VMs</b> to view data for all virtual machines on the server/host combination.
<b>Count</b>	Displays the current number of virtual machines listed in the table.

#### Fields and Data:

<b>Server</b>	The name of the server.*
<b>vSphere VM Name</b>	The name of the virtual machine.*
<b>Disk Name</b>	The name of the disk.*
<b>% Disk Usage</b>	Displays the current percentage of disk space that is being used.*
<b>Capacity (GB)</b>	Displays the total disk capacity, in gigabytes.*
<b>Disk Used (GB)</b>	Displays the total disk space currently being used.*
<b>Disk Free (GB)</b>	Displays the amount of available disk space, in gigabytes.*
<b>Host</b>	Displays the name of the host.*

**Expired** When checked, performance data for that cluster has not been received in the time specified in the **Duration** region on the RTView Configuration > (Project Name/**MISCMON-LOCAL**) > **Solution Package Configuration** > **VMWare** > **DATA STORAGE** tab.

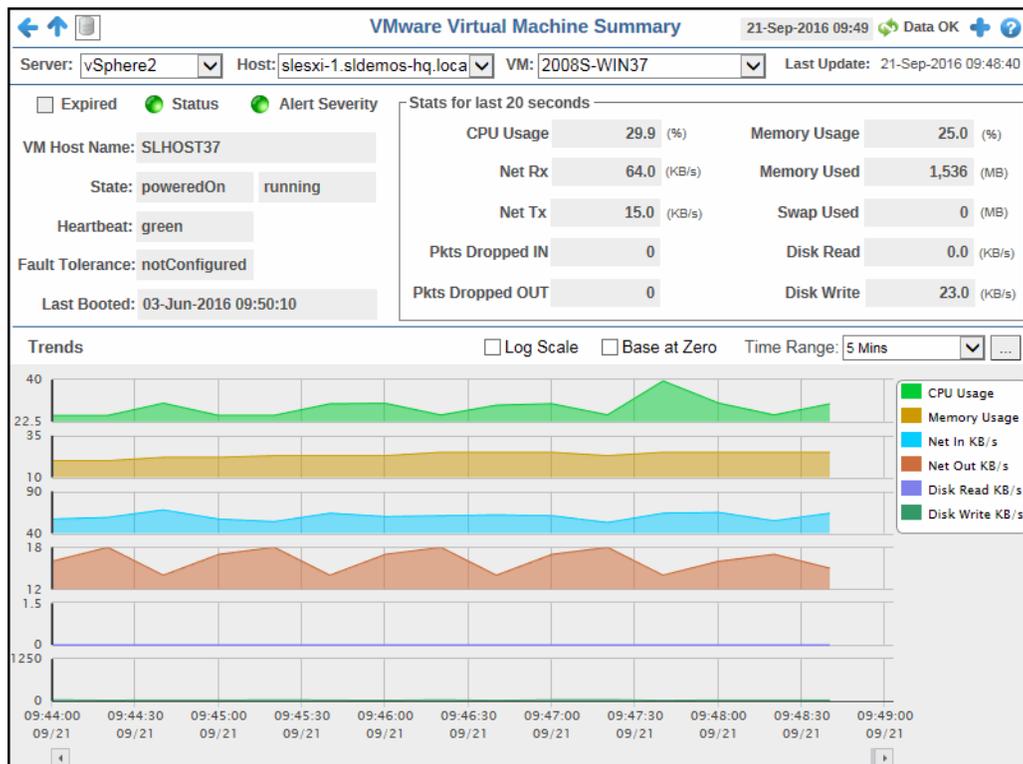
**Timestamp** The date and time the row data was last updated.

### Single VM Summary

View current and historical utilization and performance metrics for a single virtual machine. You can use this display to investigate performance issues for a particular virtual machine.

This display includes trend graphs tracing CPU and memory usage, amount of network data transmitted and received, number of incoming and outgoing packets that have been lost, and disk usage.

Choose a server, host, and virtual machine from the drop-down menus to view details for a specific virtual machine. You can use the **Time-Range** in the **Trends** region to “zoom-in” or “zoom-out” on a specific time frame in the trend graph.



**Title Bar** (possible features are):

-   Open the previous and upper display.
-  Open an instance of this display in a new window.
-  Open the online help page for this display.
-   open commonly accessed displays.
- The number of items currently in the display.

 **Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

 Open the **Alert Views - RTView Alerts Table** display.

---

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the VMWare vSphere. Refer to VMWare vsphere documentation (<http://pubs.vmware.com/vsphere-65/index.jsp#com.vmware.wssdk.apiref.doc/mo-types-landing.html>) for more information regarding these fields

---

#### Filter By:

The display might include these filtering options:

- Server** Select the server containing the virtual machine for which you want to view data.
- Host** Select the host containing the virtual machine for which you want to view data.
- VM** Select the virtual machine for which you want to view data.
- Last Update** The date and time that the data in the display was last updated.

#### Fields and Data:

- Expired** When checked, performance data for that cluster has not been received in the time specified in the **Duration** region on the RTView Configuration > (Project Name/**MISCMON-LOCAL**) > **Solution Package Configuration** > **VMWare** > **DATA STORAGE** tab.
- Status** The general health status of the virtual machine.\*
  -  Red indicates that the host is experiencing a problem.
  -  Yellow indicates that the host might have a problem.
  -  Grey indicates that the status of the host's health is unknown.
  -  Green indicates that host's status is OK.
- Alert Severity** The current severity of alerts for the virtual machine.
  -  One or more alerts exceeded their ALARM LEVEL threshold.
  -  One or more alerts exceeded their WARNING LEVEL threshold.
  -  No alert thresholds have been exceeded.
- VM Host Name** The name of the host.\*
- State** Displays whether or not the host/virtual machine is powered on.\*
- Heartbeat** Displays whether or not the virtual machine has a heartbeat.\*
  -  Red indicates that heartbeating has stopped.
  -  Grey indicates that heartbeat status is disabled.
  -  Green indicates that heartbeat status is OK.

<b>Fault Tolerance</b>	Displays whether or not fault tolerance is configured ( <b>configured/notConfigured</b> ).*	
<b>Last Booted</b>	The date and time the virtual machine was last rebooted.*	
<b>Status for last 20 seconds</b>	<b>CPU Usage</b>	The percentage of CPU used in the last 20 seconds.*
	<b>Net Rx</b>	The amount of network data received, in kilobytes per second, in the last 20 seconds.*
	<b>Net Tx</b>	The amount of network data transmitted, in kilobytes per second, in the last 20 seconds.*
	<b>Pkts Dropped IN</b>	The number of incoming packets that were dropped in the last 20 seconds.*
	<b>Pkts Dropped OUT</b>	The number of outgoing packets that were dropped in the last 20 seconds.*
	<b>Memory Usage</b>	The percentage of memory used in the last 20 seconds.*
	<b>Memory Used</b>	The amount of memory used, in megabytes, in the last 20 seconds.*
	<b>Swap Used</b>	The amount of memory swapped, in megabytes, in the last 20 seconds.*
	<b>Disk Read</b>	The amount of data read from the disk, in kilobytes per second, in the last 20 seconds.*
<b>Disk Write</b>	The amount of data written to the disk, in kilobytes per second, in the last 20 seconds.*	

### Trend Graphs

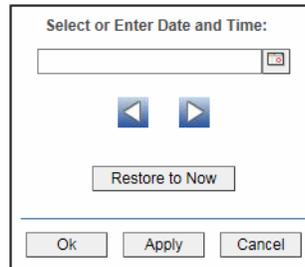
Traces the sum of process metrics for the virtual machine.

- **CPU Usage:** The percentage (%) CPU used.
- **Memory Usage:** The amount of memory used.
- **Net In KB/s:** The amount of network data received per second, in kilobytes.
- **Net Out KB/s:** The amount of network data transmitted per second, in kilobytes.
- **Disk Read KB/s:** The amount of data being read from the disk per second, in kilobytes.
- **Disk Write KB/s:** The amount of data being written to the disk per second, in kilobytes.

**Log Scale** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.

**Base at Zero** Select to use zero (0) as the Y axis minimum for all graph traces.

**Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Datstores View

The displays in this view provide a list of datstores on one or all servers, a list of all hosts mounted to a particular datstore, a list of all virtual machines hosted by a particular datstore, or data for a particular datstore. This View contains the following displays:

- **"All Datstores Table"**: View all datstores, as well as data associated with the datstores, that exist on one server or on all servers.
- **"Hosts by Datstore Table"**: View all hosts that are mounted to a particular datstore.
- **"VMs by Datstore Table"**: View all virtual machines that are hosted by a particular datstore.
- **"Single Datstore Summary"**: View metrics and trend data for a single datstore, as well as those hosts and virtual machines that are using the datstore.

## All Datastores Table

View all datastores, as well as data associated with the datastores, that exist on one server or on all servers.

Datastores							
Server	Name	Overall Status	Maintenance Mode	Type	Accessible?	Multiple Hosts?	% Utili*
qavSphere1	datastore1		normal	VMFS	<input checked="" type="checkbox"/>	<input type="checkbox"/>	
qavSphere1	datastore1 (1)		normal	VMFS	<input checked="" type="checkbox"/>	<input type="checkbox"/>	
qavSphere1	datastore-A1		normal	VMFS	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	
qavSphere1	datastore-A2		normal	VMFS	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	

**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** **Table** open commonly accessed displays.
- 6,047** The number of items currently in the display.
- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the VMWare vSphere. Refer to VMWare vsphere documentation (<http://pubs.vmware.com/vsphere-65/index.jsp#com.vmware.wssdk.apiref.doc/mo-types-landing.html>) for more information regarding these field.

**Filter By:**

The display might include these filtering options:

<b>Server</b>	Select the server for which you want to view data.
<b>Count</b>	The total number of datastores on the selected server(s), which are listed in the <b>Datastores</b> table.

**Datastores Table:**

<b>Server</b>	The name of the server.
<b>Name</b>	The name of the datastore.
<b>Overall Status</b>	The general health status of the datastore.*  Red indicates that the datastore is experiencing a problem.  Yellow indicates that the datastore might have a problem.  Grey indicates that the status of the datastore's health is unknown.  Green indicates that datastore's status is OK.
<b>Maintenance Mode</b>	Lists current maintenance mode state of the datastore (normal, inMaintenance, enteringMaintenance).*
<b>Type</b>	Lists the type of file system volume, such as VMFS or NFS.*
<b>Accessible?</b>	The connectivity status of the datastore. When checked, indicates that the datastore is accessible.*
<b>Multiple Hosts</b>	When checked, indicates that more than one host has been configured with access to the datastore.*
<b>% Utilization</b>	Lists the current space utilization percentage for the datastore.*
<b>Capacity (GB)</b>	Displays the maximum capacity of the datastore, in gigabytes.*
<b>Free Space (GB)</b>	Displays the amount of available space in the datastore, in gigabytes.*
<b>Space Uncommitted (GB)</b>	Displays the amount of total additional storage space potentially used by all virtual machines on this datastore, in gigabytes.*
<b>Expired</b>	When checked, performance data for that cluster has not been received in the time specified in the <b>Duration</b> region on the RTView Configuration > (Project Name/ <b>MISCMON-LOCAL</b> ) > <b>Solution Package Configuration</b> > <b>VMWare</b> > <b>DATA STORAGE</b> tab.
<b>Timestamp</b>	The date and time the row data was last updated.

## Hosts by Datastore Table

View all hosts that are mounted to a particular datastore.

Server	Datastore Name	Host Name	Access Mode	Accessible?	Mounted?	
qavSphere1	datastore-A1	slesxi-1.sldemos-hq.local	readWrite	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	/vmfs/volumes/511bac...
qavSphere1	datastore-A1	slesxi-2.sldemos-hq.local	readWrite	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	/vmfs/volumes/511bac...

**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.
- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the VMWare vSphere. Refer to VMWare vsphere documentation (<http://pubs.vmware.com/vsphere-65/index.jsp#com.vmware.wssdk.apiref.doc/mo-types-landing.html>) for more information regarding these fields.

### Filter By:

The display might include these filtering options:

**Server**

Select the server containing the datastore for which you want to view data, or select **All Servers**.

- Datastore** Select the datastore for which you want to view data.
- Count** The total number of hosts connecting to a datastore, which are listed in the table.

### Hosts Mounting Selected Datastore Table:

- Server** The name of the server.
- Datastore Name** The name of the datastore.
- Host Name** The name of the host.
- Access Mode** Lists the host system's access mode to the datastore (readWrite or readOnly).\*
- Accessible?** The connectivity status of the datastore. When checked, indicates that the datastore is accessible.\*
- Mounted** When checked, indicates that the datastore is mounted on the host.\*
- Mount Path** Lists the file path where the file system volume is mounted.\*
- Expired** When checked, performance data for that cluster has not been received in the time specified in the **Duration** region on the RTView Configuration > (Project Name/ **MISCOMN-LOCAL**) > **Solution Package Configuration** > **VMWare** > **DATA STORAGE** tab.
- Timestamp** The date and time the row data was last updated.

## VMs by Datastore Table

View all virtual machines that are hosted by a particular datastore.

VMware Datastore VMs - Table						
Server: All Servers		Datastore: datastore-A1		Last Update: 22-Mar-2017 11:22:00		
Count: 38		VMs Hosted by Datastore				
Server	Datastore Name	VM Name	Expired	Timestamp		
qavSphere1	datastore-A1	2008S-IT1	<input type="checkbox"/>	22-Mar-2017 11:22:0		
qavSphere1	datastore-A1	2008S-JIRA-DEV4	<input type="checkbox"/>	22-Mar-2017 11:22:0		
qavSphere1	datastore-A1	2008S-SLHOST-WIN7	<input type="checkbox"/>	22-Mar-2017 11:22:0		
qavSphere1	datastore-A1	2008S-SLHOST-WIN9	<input type="checkbox"/>	22-Mar-2017 11:22:0		
qavSphere1	datastore-A1	2008S-WIN10	<input type="checkbox"/>	22-Mar-2017 11:22:0		
qavSphere1	datastore-A1	2008S-WIN11	<input type="checkbox"/>	22-Mar-2017 11:22:0		
qavSphere1	datastore-A1	2008S-WIN37	<input type="checkbox"/>	22-Mar-2017 11:22:0		
qavSphere1	datastore-A1	2008S-WIN38	<input type="checkbox"/>	22-Mar-2017 11:22:0		
qavSphere1	datastore-A1	2008S-WIN44	<input type="checkbox"/>	22-Mar-2017 11:22:0		
qavSphere1	datastore-A1	2008S-WIN44-CLONE	<input type="checkbox"/>	22-Mar-2017 11:22:0		
qavSphere1	datastore-A1	64BIT-OL7-0	<input type="checkbox"/>	22-Mar-2017 11:22:0		
qavSphere1	datastore-A1	64BIT-OL7-14	<input type="checkbox"/>	22-Mar-2017 11:22:0		
qavSphere1	datastore-A1	64BIT-OL7-15	<input type="checkbox"/>	22-Mar-2017 11:22:0		
qavSphere1	datastore-A1	64BIT-OL7-16	<input type="checkbox"/>	22-Mar-2017 11:22:0		
qavSphere1	datastore-A1	64BIT-OL7-17	<input type="checkbox"/>	22-Mar-2017 11:22:0		
qavSphere1	datastore-A1	64BIT-OL7-6	<input type="checkbox"/>	22-Mar-2017 11:22:0		
qavSphere1	datastore-A1	64BIT-OL7-UEK4-0	<input type="checkbox"/>	22-Mar-2017 11:22:0		
qavSphere1	datastore-A1	64BIT-OL7-UEK4-DOCKER1.11-0	<input type="checkbox"/>	22-Mar-2017 11:22:0		
qavSphere1	datastore-A1	64BIT-RH6-3-0	<input type="checkbox"/>	22-Mar-2017 11:22:0		
qavSphere1	datastore-A1	64BIT-RH6-1	<input type="checkbox"/>	22-Mar-2017 11:22:0		
qavSphere1	datastore-A1	64BIT-RH6-3-1	<input type="checkbox"/>	22-Mar-2017 11:22:0		
qavSphere1	datastore-A1	64BIT-RH6-3-2	<input type="checkbox"/>	22-Mar-2017 11:22:0		
qavSphere1	datastore-A1	64BIT-RH6-3-3	<input type="checkbox"/>	22-Mar-2017 11:22:0		
qavSphere1	datastore-A1	64BIT-RH6-3-5	<input type="checkbox"/>	22-Mar-2017 11:22:0		
qavSphere1	datastore-A1	64BIT-RH6-6-8	<input type="checkbox"/>	22-Mar-2017 11:22:0		
qavSphere1	datastore-A1	64BIT-RHEL7-1	<input type="checkbox"/>	22-Mar-2017 11:22:0		
qavSphere1	datastore-A1	64BIT-SOL10-0	<input type="checkbox"/>	22-Mar-2017 11:22:0		
qavSphere1	datastore-A1	64BIT-W2008S-R2-0	<input type="checkbox"/>	22-Mar-2017 11:22:0		
qavSphere1	datastore-A1	OL7-20	<input type="checkbox"/>	22-Mar-2017 11:22:0		

**Title Bar** (possible features are):

-  Open the previous and upper display.
-  Open an instance of this display in a new window.
-  Open the online help page for this display.
- open commonly accessed displays.
- The number of items currently in the display.

 **Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

 Open the **Alert Views - RTView Alerts Table** display.

---

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the VMWare vSphere. Refer to VMWare vsphere documentation (<http://pubs.vmware.com/vsphere-65/index.jsp#com.vmware.wssdk.apiref.doc/mo-types-landing.html>) for more information regarding these field.

---

#### Filter By:

The display might include these filtering options:

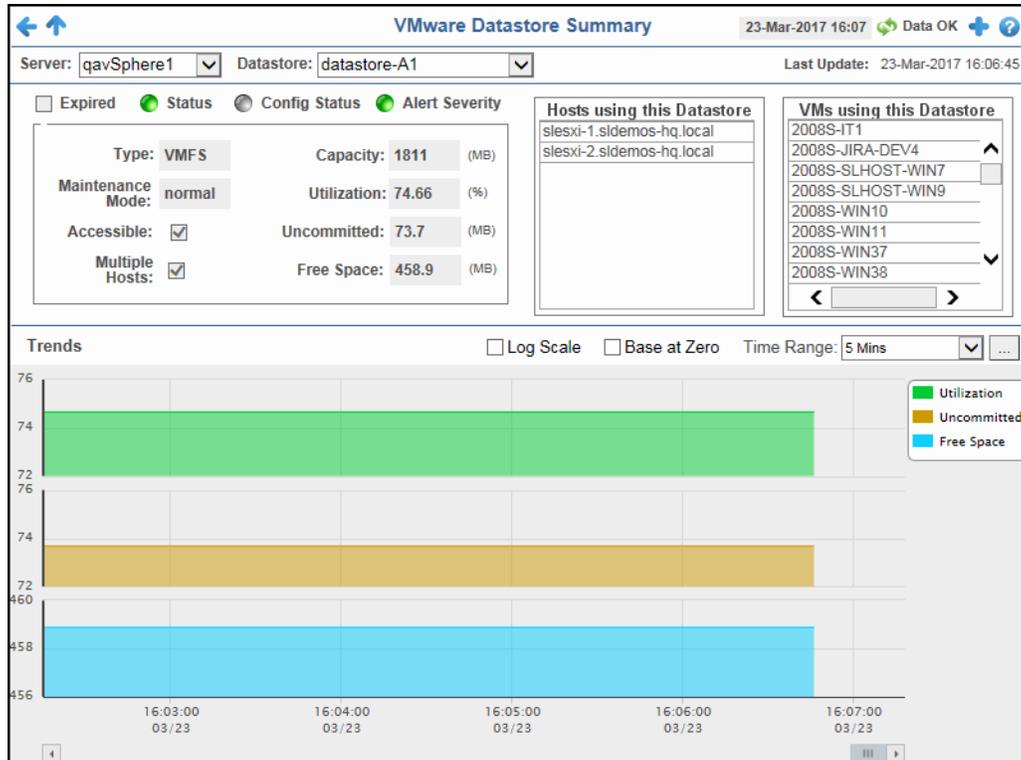
<b>Server</b>	Select the server containing the datastore for which you want to view data, or select <b>All Servers</b> .
<b>Datastore</b>	Select the datastore for which you want to view data.
<b>Count</b>	The total number of virtual machines connecting to a datastore, which are listed in the table.

#### VMs Hosted by Datastore Table:

<b>Server</b>	The name of the server.
<b>Datastore Name</b>	The name of the datastore.
<b>VM Name</b>	The name of the virtual machine hosted by the datastore.
<b>Expired</b>	When checked, performance data for that cluster has not been received in the time specified in the <b>Duration</b> region on the RTView Configuration > (Project Name/ <b>MISCMON-LOCAL</b> ) > <b>Solution Package Configuration</b> > <b>VMWare</b> > <b>DATA STORAGE</b> tab.
<b>Timestamp</b>	The date and time the row data was last updated.

## Single Datastore Summary

View metrics and trend data for a single datastore, as well as those hosts and virtual machines that are using the datastore.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- The number of items currently in the display.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

Open the **Alert Views - RTView Alerts Table** display.

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the VMWare vSphere. Refer to VMWare vsphere documentation (<http://pubs.vmware.com/vsphere-65/index.jsp#com.vmware.wssdk.apiref.doc/mo-types-landing.html>) for more information regarding these field.

**Filter By:**

The display might include these filtering options:

<b>Server</b>	Select the server containing the datastore for which you want to view data.
<b>Datastore</b>	Select the datastore for which you want to view data.
<b>Last Update</b>	The date and time the data in the display was last updated.

**Fields and Data:**

<b>Expired</b>	When checked, performance data for that cluster has not been received in the time specified in the <b>Duration</b> region on the RTView Configuration > (Project Name/ <b>MISCMON-LOCAL</b> ) > <b>Solution Package Configuration</b> > <b>VMWare</b> > <b>DATA STORAGE</b> tab.
<b>Status</b>	The general health status of the datastore.*  Red indicates that the datastore is experiencing a problem.  Yellow indicates that the datastore might have a problem.  Grey indicates that the status of the datastore's health is unknown.  Green indicates that datastore's status is OK.
<b>Config Status</b>	Indicates whether or not the system has detected a configuration issue involving the datastore.  Red indicates that a problem has been detected involving the datastore.  Yellow indicates a problem is about to occur or a transient condition has occurred.  Grey indicates that configuration status of the datastore is not being monitored.  Green indicates that no configuration issues have been detected.
<b>Alert Severity</b>	The current severity of alerts for the datastore.  One or more alerts exceeded their ALARM LEVEL threshold.  One or more alerts exceeded their WARNING LEVEL threshold.  No alert thresholds have been exceeded.
<b>Type</b>	Lists the type of file system volume, such as VMFS or NFS.*
<b>Maintenance Mode</b>	Lists current maintenance mode state of the datastore (normal, inMaintenance, enteringMaintenance).*
<b>Accessible?</b>	The connectivity status of the datastore. When checked, indicates that the datastore is accessible.*
<b>Multiple Hosts</b>	When checked, indicates that more than one host has been configured with access to the datastore.*
<b>Capacity (MB)</b>	Displays the maximum capacity of the datastore, in megabytes.*
<b>Utilization (%)</b>	Lists the current space utilization percentage for the datastore.*
<b>Uncommitted (MB)</b>	Displays the amount of total additional storage space potentially used by all virtual machines on this datastore, in megabytes.*
<b>Free Space (MB)</b>	Displays the amount of available space in the datastore, in megabytes.*

**Hosts Using this Datastore** -- Lists the hosts using the datastore.

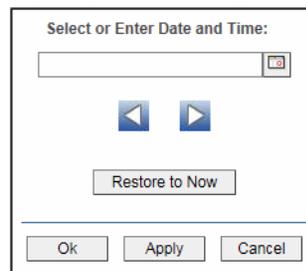
**VMs using this Datastore** -- Lists the virtual machines using the datastore.

### Trend Graphs

Traces the sum of process metrics for the virtual machine.

- **Utilization:** Traces the current space utilization percentage for the datastore.
- **Uncommitted:** Traces the amount of total additional storage space potentially used by all virtual machines on this datastore, in megabytes.
- **Free Space:** Traces the amount of available space in the datastore, in megabytes.

- Log Scale** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Base at Zero** Select to use zero (**0**) as the Y axis minimum for all graph traces.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Networks View

The display in this View lists all networks, as well as data associated with the networks, that exist on one server or on all servers. The available display in this View is:

- **"All Networks Table":** View all networks, as well as data associated with the networks, that exist on one server or on all servers.

### All Networks Table

View all networks, as well as data associated with the networks, that exist on one server or on all servers.

VMware Networks - Table 22-Mar-2017 11:37 Data OK

Server: All Servers

Count: 1

Server	Network Name	Overall Status	Accessible?	IP Pool Name	Expired	Timestamp
qavSphere1	VM Network		<input checked="" type="checkbox"/>		<input type="checkbox"/>	22-Mar-2017 11:36:...

**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu , Table open commonly accessed displays.
- 6,047 The number of items currently in the display.
- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the VMWare vSphere. Refer to VMWare vsphere documentation (<http://pubs.vmware.com/vsphere-65/index.jsp#com.vmware.wssdk.apiref.doc/mo-types-landing.html>) for more information regarding these field.

**Filter By:**

The display might include these filtering options:

- Server** Select the server for which you want to view data.
- Count** The total number of networks on the selected server(s), which are listed in the table.

**VM Networks Table:**

<b>Server</b>	The name of the server.
<b>Network Name</b>	The name of the network.
<b>Overall Status</b>	The general health status of the network.*  Red indicates that the network is experiencing a problem.  Yellow indicates that the network might have a problem.  Grey indicates that the status of the network's health is unknown.  Green indicates that network's status is OK.
<b>Accessible?</b>	The connectivity status of the virtual machine. When checked, indicates that the virtual machine is accessible.*
<b>IP Pool Name</b>	Lists the name of the IP pool that is assigned to the network.*
<b>Expired</b>	When checked, performance data for that cluster has not been received in the time specified in the <b>Duration</b> region on the RTView Configuration > (Project Name/ <b>MISCMON-LOCAL</b> ) > <b>Solution Package Configuration</b> > <b>VMWare</b> > <b>DATA STORAGE</b> tab.
<b>Timestamp</b>	The date and time the row data was last updated.

## Events/Alarms View

The displays in this View allow you to view event data and alarm data for one server or for all servers. Available displays in this View are:

- ["All Events Table"](#): View all events, as well as data associated with the events, that exist on one server or on all servers.
- ["All Alarms Table"](#): View all alarms, as well as data associated with the alarms, that exist on one server or on all servers.

## All Events Table

View all events, as well as data associated with the events, that exist on one server or on all servers.

← ↑ VMware Events - Table 22-Mar-2017 11:44 Data OK + ?

Server: All Servers ▾

Count: 245

Server	Created Time	Event ID	Chain ID	Event Class	Event Type	User Name	
qavSphere1	22-Mar-2017 08:47:22	184029	184029	UserLoginSessionEvent	info	qauser	Use
qavSphere1	22-Mar-2017 08:42:54	184028	184028	UserLoginSessionEvent	info	qauser	Use
qavSphere1	22-Mar-2017 08:41:24	184027	184027	UserLoginSessionEvent	info	qauser	Use
qavSphere1	22-Mar-2017 08:24:55	184026	184026	UserLogoutSessionEvent	info	vmwarehost1	Use
qavSphere1	22-Mar-2017 07:54:50	184025	184025	UserLoginSessionEvent	info	vmwarehost1	Use
qavSphere1	22-Mar-2017 07:54:50	184024	184024	UserLoginSessionEvent	info	qauser	Use
qavSphere1	22-Mar-2017 07:54:20	184023	184023	UserLogoutSessionEvent	info	qauser	Use
qavSphere1	22-Mar-2017 07:22:56	184022	184022	UserLogoutSessionEvent	info	qauser	Use
qavSphere1	22-Mar-2017 07:10:32	184021	184021	UserLogoutSessionEvent	info	vmwarehost1	Use
qavSphere1	22-Mar-2017 07:02:20	184020	184020	UserLogoutSessionEvent	info	qauser	Use
qavSphere1	22-Mar-2017 06:52:56	184019	184019	UserLoginSessionEvent	info	qauser	Use
qavSphere1	22-Mar-2017 06:40:29	184018	184018	UserLoginSessionEvent	info	vmwarehost1	Use
qavSphere1	22-Mar-2017 06:40:28	184017	184017	UserLoginSessionEvent	info	qauser	Use
qavSphere1	22-Mar-2017 06:38:13	184016	184016	UserLogoutSessionEvent	info	qauser	Use
qavSphere1	22-Mar-2017 06:35:38	184015	184015	UserLoginSessionEvent	info	qauser	Use
qavSphere1	22-Mar-2017 00:35:09	184014	184014	UserLogoutSessionEvent	info	qauser	Use
qavSphere1	22-Mar-2017 00:07:02	184013	184013	UserLoginSessionEvent	info	qauser	Use
qavSphere1	21-Mar-2017 18:06:58	184012	184012	UserLogoutSessionEvent	info	qauser	Use
qavSphere1	21-Mar-2017 17:44:07	184011	184011	UserLogoutSessionEvent	info	qauser	Use
qavSphere1	21-Mar-2017 17:34:10	184010	184010	UserLoginSessionEvent	info	qauser	Use
qavSphere1	21-Mar-2017 17:21:10	184009	184009	UserLogoutSessionEvent	info	qauser	Use
qavSphere1	21-Mar-2017 17:10:00	184008	184008	UserLoginSessionEvent	info	qauser	Use
qavSphere1	21-Mar-2017 17:03:30	184007	184007	AlarmStatusChangedEvent	info		Alar
qavSphere1	21-Mar-2017 17:03:09	184006	184006	AlarmStatusChangedEvent	info		Alar
qavSphere1	21-Mar-2017 17:01:49	184005	184005	AlarmStatusChangedEvent	info		Alar
qavSphere1	21-Mar-2017 15:02:39	184004	184004	UserLogoutSessionEvent	info	qauser	Use
qavSphere1	21-Mar-2017 14:53:56	184003	184003	UserLogoutSessionEvent	info	vmwarehost1	Use

Page 1 of 2 1 - 200 of 245 items

**Title Bar (possible features are):**

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu ▾, Table open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the VMWare vSphere. Refer to VMWare vsphere documentation (<http://pubs.vmware.com/vsphere-65/index.jsp#com.vmware.wssdk.apiref.doc/mo-types-landing.html>) for more information regarding these field.

**Filter By:**

The display might include these filtering options:

<b>Server</b>	Select the server for which you want to view data.
<b>Count</b>	The total number of events on the selected server(s), which are listed in the <b>Events</b> table.

**Events Table:**

<b>Server</b>	The name of the server.
<b>Created Time</b>	The date and time the event was created.*
<b>Event ID</b>	The ID of the event.*
<b>Chain ID</b>	The parent or group ID.*
<b>Event Class</b>	The type of event class.*
<b>Event Type</b>	The type of event.*
<b>User Name</b>	The user who caused the event.*
<b>Message Text</b>	A formatted text message describing the event.*
<b>Host</b>	The host object of the event.*
<b>Virtual Machine</b>	The event's virtual machine.*
<b>Compute Resource</b>	The event's compute resource.*
<b>Datacenter</b>	The event's datacenter.*
<b>Datastore</b>	The event's datastore.*
<b>Distributed Virtual Switch</b>	The event's DistributedVirtualSwitch.*
<b>Network</b>	The network associated with the event.*
<b>Expired</b>	When checked, performance data for that cluster has not been received in the time specified in the <b>Duration</b> region on the RTView Configuration > (Project Name/ <b>MISCMON-LOCAL</b> ) > <b>Solution Package Configuration</b> > <b>VMWare</b> > <b>DATA STORAGE</b> tab.
<b>Timestamp</b>	The date and time the row data was last updated.

## All Alarms Table

View all alarms, as well as data associated with the alarms, that exist on one server or on all servers.

Alarm Time	Server	Managed Object	Overall Status	Alarm Name
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	●	Host connection and power state
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	●	Timed out starting Secondary VM
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	●	No compatible host for Secondary VM
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	●	Host processor status
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	●	Host memory status
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	●	Host hardware fan status
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	●	Host hardware voltage
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	●	Host hardware temperature status
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	●	Host hardware power status
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	●	Host hardware system board status
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	●	Host battery status
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	●	Virtual Machine Fault Tolerance vLockStep int
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	●	Status of other host hardware objects
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	●	Host storage status
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	●	Host error
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	●	Virtual machine error
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	●	Host connection failure
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	●	Cannot connect to storage
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	●	Migration error
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	●	Exit standby error
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	●	License error
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	●	Health status changed alarm
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	●	Host cpu usage
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	●	Virtual machine Fault Tolerance state change
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	●	Network connectivity lost
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	●	Network uplink redundancy lost
24-Jan-2017 10:00:50	qavSphere1	New Datacenter	●	Network uplink redundancy degraded

**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** **Table** open commonly accessed displays.
- 6,047** The number of items currently in the display.
- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the VMWare vSphere. Refer to VMWare vsphere documentation (<http://pubs.vmware.com/vsphere-65/index.jsp#com.vmware.wssdk.apiref.doc/mo-types-landing.html>) for more information regarding these field.

**Filter By:**

The display might include these filtering options:

<b>Server</b>	Select the server for which you want to view data.
<b>Count</b>	The total number of datastores on the selected server(s), which are listed in the <b>Datastores</b> table.

**Alarms Table:**

<b>AlarmTime</b>	The date and time of the alarm.*
<b>Server</b>	The name of the server in which the alarm occurred.*
<b>Managed Object</b>	The name of the managed object.*
<b>Overall Status</b>	The general health status of the alarm.*  Red indicates that the alarm is experiencing a problem.  Yellow indicates that the alarm might have a problem.  Grey indicates that the status of the alarm's health is unknown.  Green indicates that alarm's status is OK.
<b>Alarm Name</b>	The name of the alarm.*
<b>Acknowledged</b>	Lists whether or not the alarm has been acknowledged.*
<b>Acknowledged By</b>	Lists the user who acknowledged the alarm.*
<b>Acknowledged Time</b>	Lists the date and time when the alarm was acknowledged.*
<b>Description</b>	The description of the alarm.*
<b>Timestamp</b>	The date and time the row data was last updated.

---

## VMware vCenter - HTML

This section describes the HTML version of the Solution Package for VMware vCenter which features an overview display, "[VMware Overview Display](#)" (shown below), and the following displays which can be found under **Components** tab > **Hosts** > **VMware**. For additional details, see vendor documentation.

This section includes the following Views/Displays:

- [“VMware Overview Display”](#): Selecting VMware from the left/navigation menu opens this display, which provides a high level overview of your virtual machines and hosts.
- [“VMware Hosts View - HTML”](#): View performance and utilization data for hosts running on one or all clusters, view utilization data for a specific host running virtual machines, view a list of components contained on a selected host, and view physical and virtual network adapters located on a particular host.
- [“VMware Machines View - HTML”](#): View current and historical data for your virtual machines.

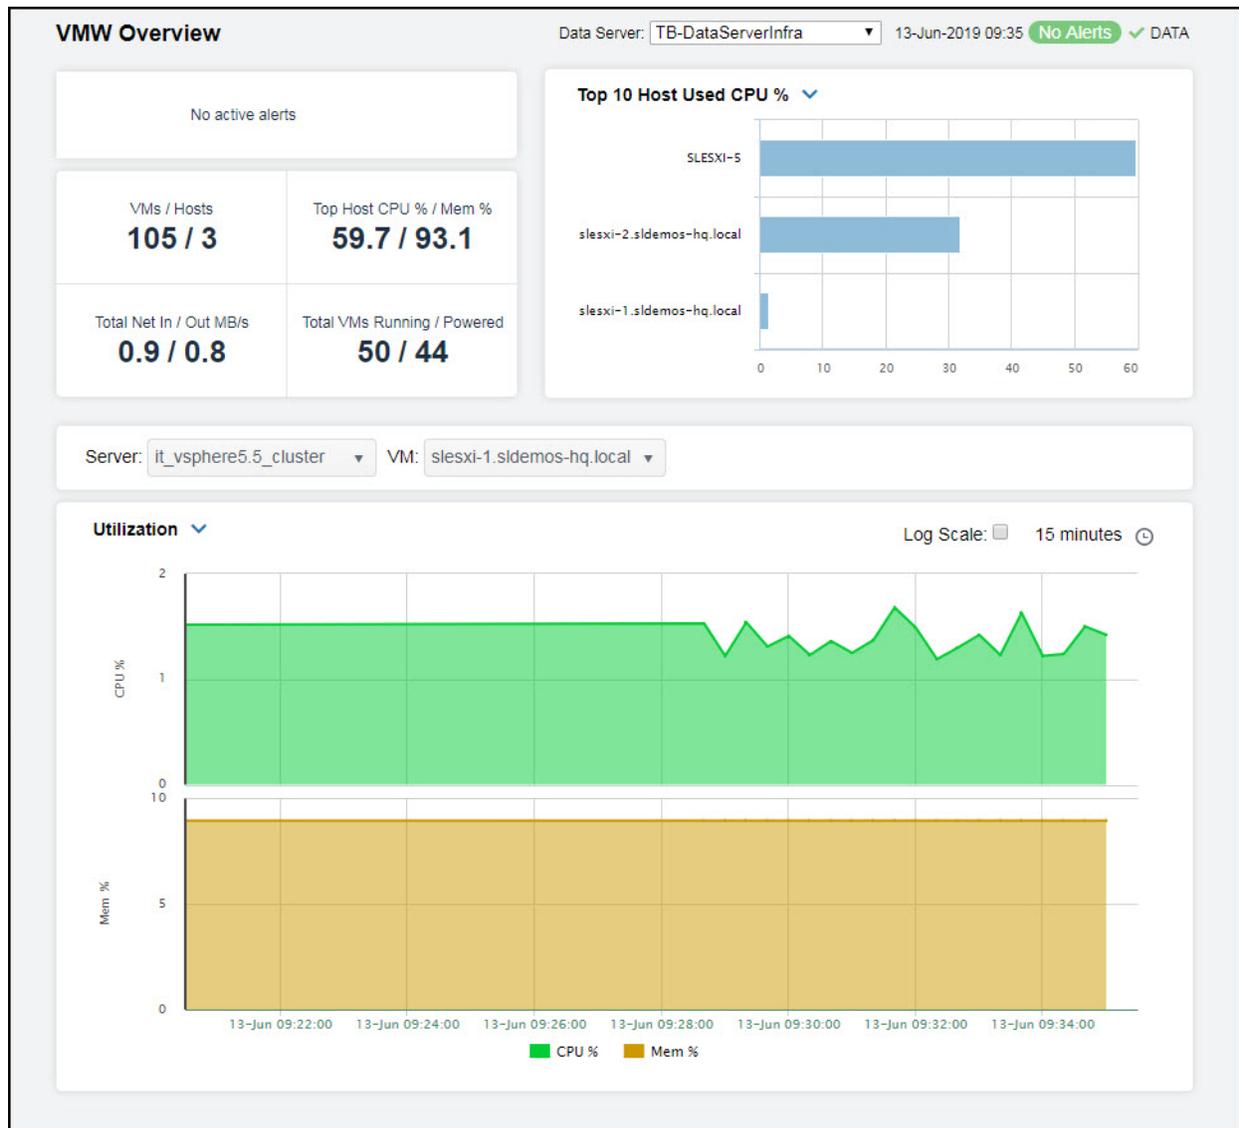
## VMware Overview Display

The **VMware Overview** is the top-level display for the TIBCO VMware Solution Package, which provides a good starting point for immediately getting the status of all your virtual machines and hosts on your Data Server. You can select the RTView DataServer for which you want to see data and easily view the current data for that DataServer including:

- The total number of active alerts for the selected DataServer, including the total number of critical and warning alerts.
- The number of virtual machines across all hosts and the total number of hosts.
- The maximum percentage of CPU utilization and the maximum percentage of memory utilization across all hosts on your connected DataServer.
- The total inbound traffic and total outbound traffic across all hosts on your connected DataServer.
- The total number of virtual machines running and the total number of virtual machines powered on across all hosts on your connected DataServer.
- A visual list of the top 10 hosts with the highest used CPU percentage, hosts with the highest memory used percentage, virtual machines with the most incoming network traffic (in kilobytes), virtual machines with the most outgoing network traffic (in kilobytes), virtual machines with the most disk reads, and the virtual machines with the most disk writes on your connected DataServer.

You can hover over each region in the upper half of the Overview to see more detail. You can also drill down to see even more detail by clicking on each respective region in the Overview. For example, clicking on the alerts in the **CRITICAL** and **WARNING** alerts region opens the **Alerts Table by Components** display.

The bottom half of the display provides several trend graphs for a particular virtual machine. You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.



## VMware Hosts View - HTML

These displays present performance and utilization data for hosts running on one or all clusters and present utilization data for a specific host running virtual machines. Clicking **VMware Hosts** from the left/navigation menu opens the [“VMware Hosts Table - HTML”](#) display, which shows a tabular view of the utilization data for all hosts running on one or on all clusters. The option available under **VMware Hosts** is:

- **Single VMware Host:** Opens the [“VMware Host Summary - HTML”](#), which shows utilization data for a specific host running virtual machines.

## VMware Hosts Table - HTML

View the utilization data for all hosts running on one cluster or on all clusters. Each row in the table contains data for a particular host. Click a column header to sort column data in ascending or descending order. Double-click on a table row to drill-down to the ["VMware Host Summary - HTML"](#) display and view metrics for that particular host.

**VMware Hosts Table**
13-Jun-2019 11:05 No Alerts DATA

Server: it\_vsphere5.5\_cluster

Hosts: **2**

**Hosts**

Server	Host Name	Alert Level	Alert Count	Overall Status	Connection State
it_vsphere5.5_clust	slesxi-2.sldemos-hq.local	✓		green	
it_vsphere5.5_clust	slesxi-1.sldemos-hq.local	✓		green	

## VMware Host Summary - HTML

Clicking **Single VMware Host** in the left/navigation menu opens the **VMware Host Summary** display, which allows you to view the number of virtual machines running on a particular host, the most recent utilization data for the host, and the trend data for the host over a specified time range. Clicking on the information boxes at the top of the display takes you to the "[VMware Hosts Table - HTML](#)" display, where you can view additional host data.

In the trend graph region, you can select from **Utilization**, **Disk**, **Network**, or **Network Loss**. **Utilization** traces CPU and Memory usage percentage over a specified period of time. **Disk** traces disk reads, disk writes, CPU used percentage, and memory used percentage over a specified period of time. **Network** traces net transactions received and transmitted, as well as CPU and Memory usage percentage over a specified period of time. **Network Loss** traces the number of packets received and transmitted over a specified period of time.

Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display.

**VMware Host Summary** ▼
13-Jun-2019 11:03 No Alerts ✓ DATA

Server: it\_vsphere5.5\_cluster ▼
Host: slesxi-1.sldemos-hq.local ▼

VMs Hosted  
1

CPU %  
1.3

Used Memory %  
9.0

Network Rx / Tx KB/s  
47.0 / 9.0

Dropped Packets Rx / Tx  
0.0 / 0.0

Disk Reads / Writes KB/s  
12.0 / 17.0

**Utilization** ▼
Log Scale:  15 minutes ⌵

■ CPU %
■ Mem %

Cluster Name: <b>SL CORP HA CLUSTER</b> <span style="color: blue; text-decoration: underline;">Critical/Warning: 0/0</span>		
Connection State:	Power State:	VMs Running: <b>15</b>
VMs Powered On: <b>15</b>	Dropped Rx Packets %: <b>0.0</b>	Dropped Tx Packets %: <b>0.0</b>
Avg Swap Used MB: <b>0.0</b>	Error Rx Packets %: <b>0.0</b>	Error Tx Packets %: <b>0.0</b>
CPU Threads: <b>24</b>		
Last Update: <b>13-Jun-2019 11:02:52</b>		

## VMware Machines View - HTML

These displays present current and historical data for your virtual machines. Clicking **VMware Machines** from the left/navigation menu opens the ["VMware Virtual Machines Table - HTML"](#) display, which shows all available data for each virtual machine by server. The option available under **VMware Machines** is:

- **Single VMware VM:** Opens the ["VMware Virtual Machine Summary - HTML"](#), which shows current and historical utilization and performance metrics for a single virtual machine.

## VMware Virtual Machines Table - HTML

Use this display to view all available data for each virtual machine by server.

Each row in the table lists the details for a virtual machine. Choose a server from the drop-down menus to view all virtual machines running on that server. You can click a column header to sort column data in numerical or alphabetical order.

Drill-down and investigate by double-clicking a row to view details for the selected virtual machine in the ["VMware Virtual Machine Summary - HTML"](#) display.

**VMware Virtual Machines Table** 13-Jun-2019 11:28 No Alerts ✓ DATA

Server:

VMs: 15

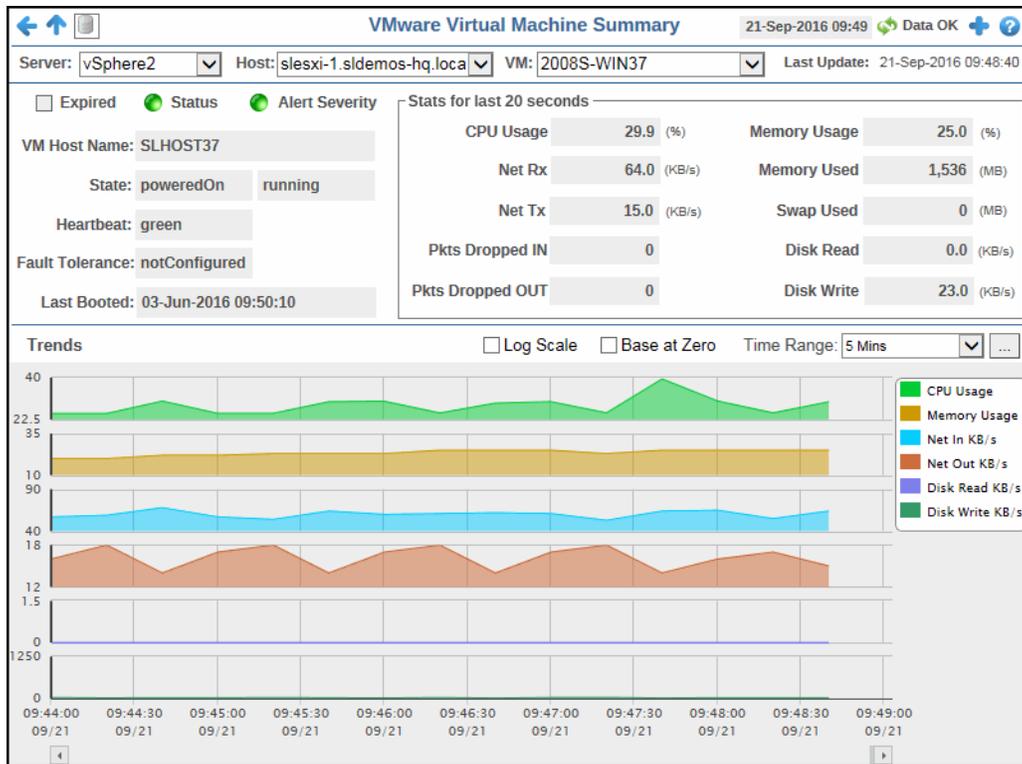
Server	Name	Alert Level	Alert Count	Heartbeat Status	Power State	
it_vsphere5.5_clust	2008S-WIN18	✓		green	poweredOn	ru ▲
it_vsphere5.5_clust	QAWIN10	✓		green	poweredOn	ru
it_vsphere5.5_clust	CENTOS7-3	✓		green	poweredOn	ru
it_vsphere5.5_clust	2008S-SLHOST-WIN8-CLC	✓		green	poweredOn	ru
it_vsphere5.5_clust	2008S-WIN13	✓		green	poweredOn	ru
it_vsphere5.5_clust	CENTOS7-4	✓		green	poweredOn	ru
it_vsphere5.5_clust	QAWIN2	✓		green	poweredOn	ru
it_vsphere5.5_clust	CENTOS7-1	✓		green	poweredOn	ru
it_vsphere5.5_clust	2008S-WIN44	✓		green	poweredOn	ru
it_vsphere5.5_clust	2008S-WIN17	✓		green	poweredOn	ru ▼

## VMware Virtual Machine Summary - HTML

Clicking **Single VMware VM** in the left/navigation menu opens the **VMware Virtual Machine Summary** display, which allows you to investigate performance issues for a particular virtual machine. Clicking on the information boxes at the top of the display takes you to the [“VMware Virtual Machines Table - HTML”](#) display, where you can view additional virtual machine data.

In the trend graph region, you can select from **Utilization**, **Disk**, **Network**, or **Network Loss**. **Utilization** traces CPU and Memory usage percentage over a specified period of time. **Disk** traces rate of disk reads and disk writes, CPU usage percentage, and memory usage percentage over a specified period of time. **Network** traces the rate of net transactions received and transmitted, as well as CPU and Memory usage percentage over a specified period of time. **Network Loss** traces the number of dropped incoming and outgoing packets over a specified period of time.

Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display.





## CHAPTER 5 RTView DataServer for Kafka

The RTView DataServer for Kafka provides a way to create connections and modify default configuration settings for solution packages and sends collected data to RTViewCentral. RTViewCentral contains the displays associated with the RTView DataServer for Kafka which you use to monitor your Kafka components. Both the display server user interface and the HTML user interface are described here.

The RTView *DataCollector* for Kafka is also available for use with the RTView DataServer for Kafka. RTView DataCollector for Kafka is used for collecting data and sending it to one or more RTView DataServers. The RTView DataCollector for Kafka is useful if you need to distribute data collection.

For an overview and details about configuring RTView Enterprise, including RTViewCentral, RTView DataServers, RTView DataCollectors and solution packages, see the *RTView Enterprise Configuration Guide*.

RTViewCentral contains the following displays that will be populated with data collected via the RTView DataServer for Kafka:

- "Kafka"
- "Kafka - HTML"

**Note:** This document assumes familiarity with the products monitored. For additional details, refer to vendor documentation.

---

### Kafka

The following Apache Kafka Views (and their associated displays) can be found under **Components** tab > **Middleware** > **Kafka**.

The following views are available:

- **"Kafka Clusters View"**: The displays in this View allow you to view metrics for all Kafka clusters and view the performance metrics for all servers on a particular cluster.
- **"Kafka Topics View"**: This displays in this View allow you to view metrics for all topics for a particular broker in heatmap/table format, view current and trend data for a single topic, view the metrics for all topics on a particular cluster, and view metrics for a particular topic on a particular cluster.
- **"Kafka Brokers View"**: The displays in this View allow you to view the current and historical metrics for all brokers in heatmap/table formats, view various metrics for a particular broker, and view metrics and trend data for a particular broker.
- **"Kafka Zookeepers View"**: The displays in this View allow you to view the current and historical metrics for all zookeepers in a particular cluster in heatmap/tabular format, or view current and historical metrics and trend data for a single zookeeper.
- **"Kafka Producers View"**: The displays in this View allow you to view the current and historical metrics for all producers in a particular cluster in heatmap/tabular format, or view current and historical metrics and trend data for a single producer.
- **"Kafka Consumers View"**: The displays in this View allow you to view the current and historical metrics for all consumers in a particular cluster in heatmap/tabular format, or view current and historical metrics and trend data for a single consumers.

## Kafka Clusters View

These displays allow you to view metrics for all Kafka clusters and view the performance metrics for all servers on a particular cluster. Displays in this View are:

- **"All Clusters Table"**: A tabular view of all clusters and their associated metrics.
- **"Cluster Performance"**: This display allows you to view performance metrics for all servers on a particular cluster.

## All Clusters Table

The table in this display provides a view of all of your clusters and their associated metric data including maximum alert severity, alert count, and the current value of each gathered metric. You can click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by clicking a row to view details for the selected cluster in the RT Cluster Performance display.

Cluster Name	MaxSeverity	AlertCount	# Brokers Monitored	# Brokers Running	# Active Controllers	Offli
SL-QA-Cluster1	●	0	3	2	1	
SL-QA-Cluster2	●	0	3	2	1	

### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

---

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the selected cluster. Refer to KAFKA documentation for more information regarding these fields.

---

### Fields and Data:

#### Count

Lists the number of brokers found as a result of the cluster that was selected and displayed in the Kafka Brokers table.

**Kafka Clusters Table:**

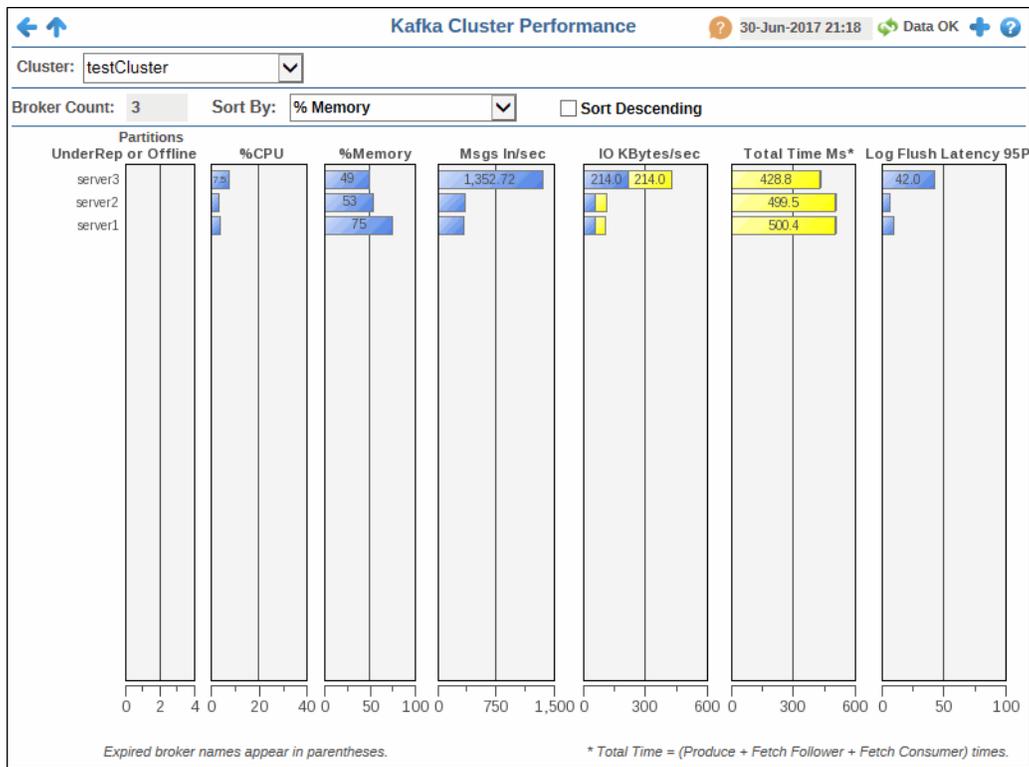
<b>Cluster Name</b>	The name of the cluster.*
<b>MaxSeverity</b>	The current highest alert severity for any of clusters.  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
<b>Alert Count</b>	The total number of alerts for the host.
<b># Brokers Monitored</b>	The current number of brokers being monitored for the cluster.*
<b># Brokers Running</b>	The number of brokers currently running on the cluster.*
<b># Active Controllers</b>	The number of active controllers on the cluster.*
<b>Offline Partitions</b>	The number of partitions without an active leader on the cluster.*
<b>Under-Replicated Partitions</b>	The number of partition replicas that are out of sync (total number of replicas minus the total number of in-sync replicas) on all brokers on the cluster.*
<b>% Max Deviation in Partition Count</b>	The percentage of maximum deviation in partition count.*
<b>% Max Deviation in Leader Count</b>	The percentage of maximum deviation in leader count.*
<b># Zookeepers</b>	The number of zookeepers on the cluster.*
<b># Zookeeper Connections</b>	The number of connections on the zookeepers in the cluster.*
<b># Zookeeper Outstanding Reqs</b>	The number of outstanding requests on the zookeepers in the cluster.
<b># Zookeeper Pkts Recvd</b>	The number of packets received on the zookeepers in the cluster.*
<b># Zookeeper Pkts Sent</b>	The number of packets sent by the zookeepers in the cluster.*
<b># Consumers</b>	The number of consumers on the cluster.*
<b>Bytes Consumed Rate</b>	The rate of bytes being consumed by the consumers.
<b>Records Consumed Rate</b>	The rate of records being consumed by the consumers.
<b># Producers</b>	The number of producers on the cluster.
<b>Producer In Byte Rate</b>	The rate of incoming bytes for the producers.*

**Producer Out Byte Rate** The rate of outgoing bytes for the producers.\*

**Producer Record Send Rate** The rate of records being sent for the producers.\*

## Cluster Performance

This display provides a view of the current metrics for the brokers contained in a selected cluster.



**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Filter By:**

The display might include these filtering options:

<b>Cluster</b>	Select the cluster for which you want to show data in the display.
<b>Sort By</b>	Select the metric by which you want to sort the data in the display. When using this option with the <b>Sort Descending</b> toggle, the brokers (servers) will be sorted in ascending or descending order using the option you select from this drop down. For example, if you select <b>MsgsInPerSec</b> from this drop down and select the <b>Sort Descending</b> toggle, the servers listed in the display will be sorted so that the server with the most <b>MsgsInPerSec</b> will be listed at the top followed by the server with the next most <b>MsgsInPerSec</b> , and so on.
<b>Sort Descending</b>	When toggled on, the servers listed in the display are sorted in descending order based on the selected metric in the <b>Sort By</b> drop down. When toggled off, the servers are listed in ascending order.

**Fields and Data:**

<b>Broker Count</b>	The number of brokers contained in the selected cluster.
<b>Partitions Under Rep or Offline</b>	Lists the number of partitions that are under-replicated or offline on each broker in the cluster.
<b>% CPU</b>	Lists the percentage of CPU used by the broker.
<b>% Memory</b>	Lists the percentage of memory used by the broker.
<b>Msgs In/sec</b>	Lists the rate of incoming messages (per second) for each broker in the cluster.
<b>IO KBytes/sec</b>	Lists the rate of incoming kilobytes (per second) for each broker in the cluster.
<b>Total Time Ms</b>	Lists the total time taken to service a request.*
<b>Log Flush Latency 95P</b>	Lists the 95th percentile value for the log flush latency for each broker on the cluster.

## Kafka Topics View

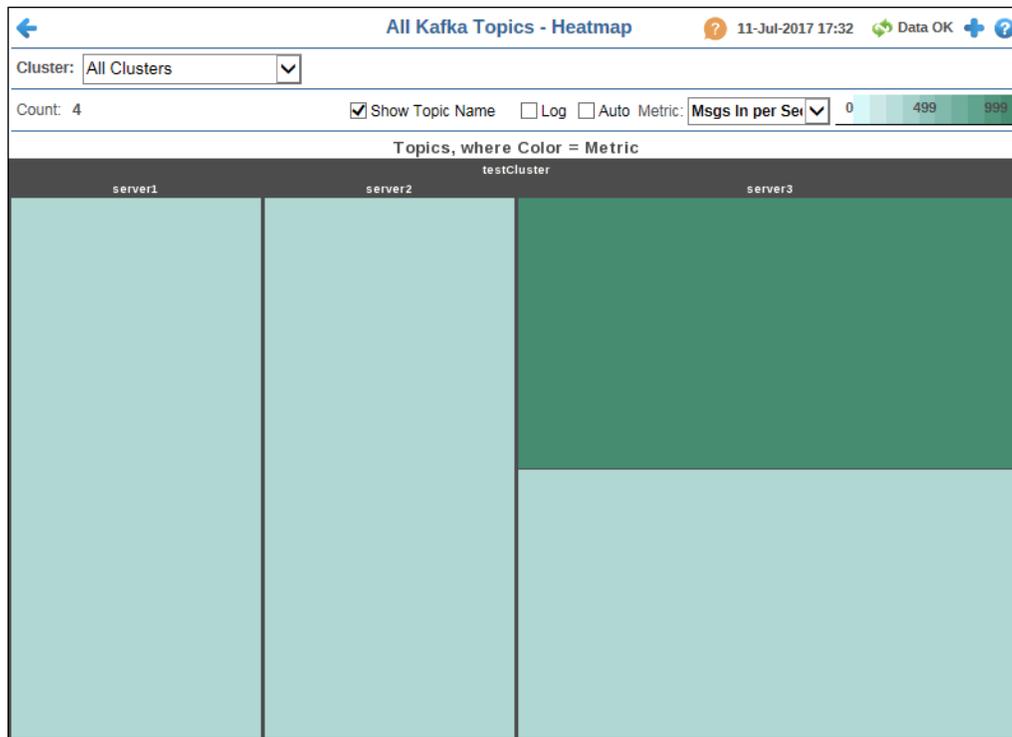
These displays allow you to view metrics for all Kafka topics on a particular topic in heatmap/table format, view the performance metrics for a single topic on a particular broker, view the metrics for all topics on a particular cluster, and view metrics for a particular topic on a particular cluster. Displays in this View are:

- ["All Topics Heatmap"](#): A heatmap view of all topics for a particular cluster.
- ["All Topics Table"](#): A tabular view of all topics for a particular cluster.
- ["Single Topic Summary"](#): This display allows you to view current metrics and trend data for a single topic.
- ["All Topics for Cluster"](#): This display allows you to view performance metrics for all topics on a particular cluster.
- ["Single Topic for Cluster"](#): This display allows you to view performance metrics for a particular topic on a particular cluster.

### All Topics Heatmap

This heatmap provides an easy-to-view interface that allows you to quickly identify the current status of each of your topics for each available metric. You can view the topics in the heatmap based on the following metrics: the rate of incoming messages, the rate of incoming bytes, the rate of outgoing bytes, the rate of rejected bytes, the rate of total fetch requests, the rate of failed fetch requests, the rate of total produce requests, and the rate of failed produce requests. By default, this display shows the heatmap based on the **Msgs In per Sec** metric.

You can use the **Show Topic Name** check-box  to include or exclude labels in the heatmap, and you can mouse over a rectangle to see additional metrics for a topic. Clicking one of the rectangles in the heatmap opens the ["Single Topic Summary"](#) display, which allows you to see additional details for the topic metrics for the selected topic.



#### Title Bar (possible features are):

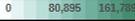
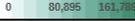
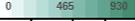
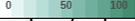
- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

#### Fields and Data:

- Cluster** Select the cluster for which you want to view data.
- Count** Lists the number of topics displayed in the heatmap.
- Show Topic Name** Select this check box to display the names of the topics at the top of each rectangle in the heatmap.
- Log** Select this check box to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Auto** Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value.  
**Note:** Some metrics auto-scale automatically, even when **Auto** is not selected.
- Metric** Choose a metric to view in the display.

<b>Msgs In per Sec</b>	The rate of incoming messages (per second). The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum number of incoming messages per second. The middle value in the gradient bar indicates the middle value of the range.
<b>Bytes In per Sec</b>	The rate of incoming bytes (per second). The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum number of incoming bytes per second. The middle value in the gradient bar indicates the middle value of the range.
<b>Bytes Out per Sec</b>	The rate of outgoing bytes (per second). The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum number of outgoing bytes per second. The middle value in the gradient bar indicates the middle value of the range.
<b>Bytes Rejected per Sec</b>	The rate of bytes being rejected (per second). The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum number of bytes rejected per second. The middle value in the gradient bar indicates the middle value of the range.
<b>Total Fetch Requests per Sec</b>	The rate of fetch requests (per second). The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum number of fetch requests per second. The middle value in the gradient bar indicates the middle value of the range.
<b>Failed Fetch Requests per Sec</b>	The rate of failed fetch requests (per second). The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum number of failed fetch requests per second. The middle value in the gradient bar indicates the middle value of the range.
<b>Total Produce Requests per Sec</b>	The rate of total producer requests (per second). The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum number of produce requests per second. The middle value in the gradient bar indicates the middle value of the range.
<b>Failed Produce Requests per Sec</b>	The rate of failed producer requests (per second). The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum number of failed produce requests per second. The middle value changes accordingly to indicate the color of the middle value of the range.

## All Topics Table

The table in this display provides a view of all of your topics for a particular cluster and their associated metric data. You can click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by clicking a row to view details for the selected topic in the ["Single Topic Summary"](#) display.

Cluster Name	Topic	BytesInPerSec	BytesOutPerSec	BytesRejectedPerSec	FailedFetchRequestsPerSec
SL-QA-Cluster1	my-replicated-topic	161,756.84	161,757.25	0.00	
SL-QA-Cluster1	my-replicated-topic1	161,756.80	165,651.96	0.00	

**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the selected cluster. Refer to KAFKA documentation for more information regarding these fields.

**Filter By**

**Cluster** Select the cluster for which you want to view data.

**Kafka Topics Table:**

**Count** The total number of topics listed in the table.

**<Rate Drop  
Down List>**

Select the option for which you want to view data:

**MeanRate** -- Select this option to view the average rate for each metric for the topics in the display.

**One Minute**-- Select this option to view the rate of incoming messages (per second) averaged over a one minute period for each metric for the topics in the display.

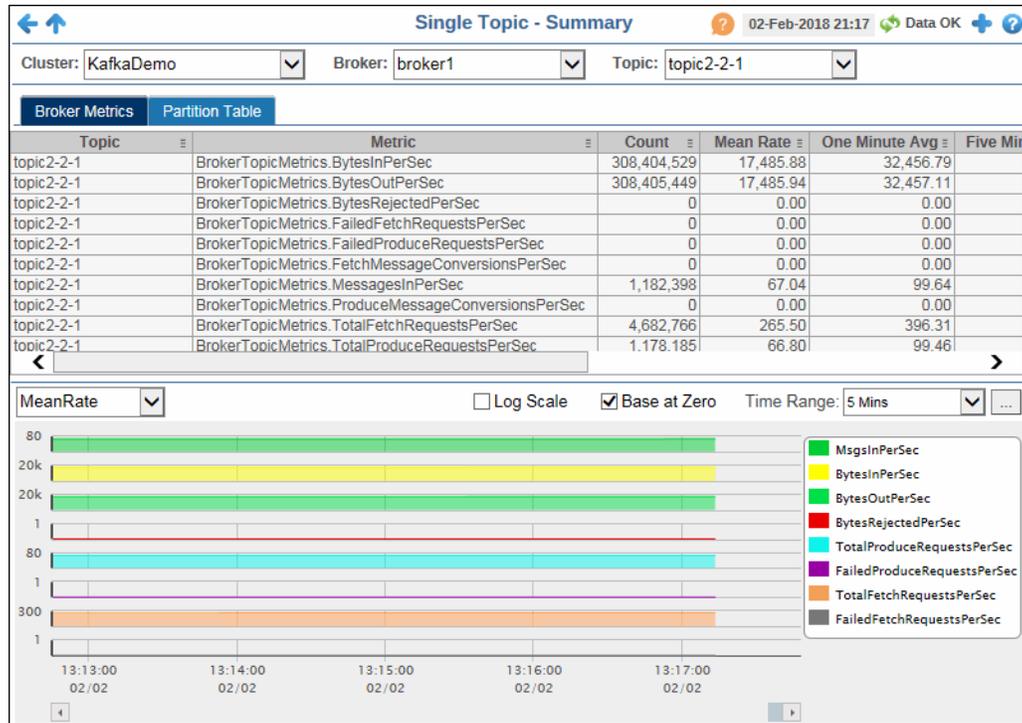
**Five Minute** -- Select this option to view the rate of incoming messages (per second) averaged over a five minute period for each metric for the topics in the display.

**Fifteen Minute** -- Select this option to view the rate of incoming messages (per second) averaged over a fifteen minute period for each metric for the topics in the display.

<b>Cluster Name</b>	Lists the name of the cluster.
<b>Topic</b>	Lists the name of the topic.
<b>Bytes In Per Sec</b>	The rate of incoming bytes.
<b>Bytes Out Per Sec</b>	The rate of outgoing bytes.
<b>Bytes Rejected Per Sec</b>	The rate of rejected bytes.
<b>Failed Fetch Requests Per Sec</b>	The rate of failed fetch requests.
<b>Failed Produce Requests Per Sec</b>	The rate of failed produce requests.
<b>Messages In Per Sec</b>	The rate of incoming messages.
<b>Total Fetch Requests Per Sec</b>	The rate of total fetch requests.
<b>Total Produce Requests Per Sec</b>	The rate of total produce requests.
<b>Replicas</b>	The total number of replicas associated with the topic.
<b>Replicas Out Of Sync</b>	The number of replicas that are out of sync for the topic.

## Single Topic Summary

This display provides a view of the current metrics and trend data for a single topic. Selecting the **Cluster/Broker** combination populates the **Topic** drop down list, from which you can select the topic for which you want to view metric and trend data.



### Title Bar (possible features are):

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu, Table open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the selected cluster. Refer to KAFKA documentation for more information regarding these fields.

### Filter By:

The display might include these filtering options:

- Cluster** Select the cluster for which you want to see data.
- Broker** Select the broker for which you want to see data.

**Topic** Select the topic for which you want to see data.

### Broker Metrics Tab

Topic	Metric	Count	Mean Rate	One Minute Avg	Five Min
topic2-2-1	BrokerTopicMetrics.BytesInPerSec	312,980,410	17,605.18	32,577.67	
topic2-2-1	BrokerTopicMetrics.BytesOutPerSec	312,981,330	17,605.24	32,577.70	
topic2-2-1	BrokerTopicMetrics.BytesRejectedPerSec	0	0.00	0.00	
topic2-2-1	BrokerTopicMetrics.FailedFetchRequestsPerSec	0	0.00	0.00	
topic2-2-1	BrokerTopicMetrics.FailedProduceRequestsPerSec	0	0.00	0.00	
topic2-2-1	BrokerTopicMetrics.FetchMessageConversionsPerSec	0	0.00	0.00	
topic2-2-1	BrokerTopicMetrics.MessagesInPerSec	1,196,438	67.30	99.96	
topic2-2-1	BrokerTopicMetrics.ProduceMessageConversionsPerSec	0	0.00	0.00	
topic2-2-1	BrokerTopicMetrics.TotalFetchRequestsPerSec	4,738,723	266.55	398.50	
topic2-2-1	BrokerTopicMetrics.TotalProduceRequestsPerSec	1,192,206	67.06	99.81	

**Topic** The name of the topic.

**Metric** The name of the metric.

**Count** The total number of the particular metric.

**Mean Rate** The mean rate of the metric.

**One Minute Avg** The rate of incoming messages (per second) averaged over a one minute period for the metric, based on the **Rate Units**.

**Five Minute Avg** The rate of incoming messages (per second) averaged over a five minute period for the metric, based on the **Rate Units**.

**Fifteen Minute Avg** The rate of incoming messages (per second) averaged over a fifteen minute period for the metric, based on the **Rate Units**.

**Event Type** The event type for the topic metric.

**Rate Units** The unit of measure used to calculate the **Mean Rate**, the **One Minute Avg**, the **Five Minute Avg**, and the **Fifteen Minute Avg**.

**Expired** When checked, performance data in the row has not been received within the time specified (in seconds) in the **Expire Time** field in the RTView Configuration Application > (**KAFKAMON-LOCAL/Project Name**) > **Solution Package Configuration** > **Apache Kafka** > **DATA STORAGE** > **Duration** > **Expire Time** property. The RTView Configuration Application > (**KAFKAMON-LOCAL/Project Name**) > **Solution Package Configuration** > **Apache Kafka** > **DATA Storage** > **Duration** > **Delete Time** property allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.

For example, if **Expire Time** was set to 120 and **Delete Time** was set to 3600, then the **Expired** check box would be checked after 120 seconds and the row would be removed from the table after 3600 seconds.

**time\_stamp** The date and time the row data was last updated.

**Rate Trends**

Select the option for which you want to view data:

**MeanRate** -- Select this option to view the average rate for each metric in the trend graph.

**One Minute**-- Select this option to view the rate averaged over a one minute period for each metric in the trend graph.

**Five Minute** -- Select this option to view the rate averaged over a five minute period for each metric in the trend graph.

**Fifteen Minute** -- Select this option to view the rate averaged over a fifteen minute period for each metric in the trend graph.

Traces the following:

**MsgsInPerSec** -- traces the selected rate of incoming messages.

**BytesInPerSec** -- traces the selected rate of incoming bytes.

**BytesOutPerSec** -- traces the selected rate of outgoing bytes.

**BytesRejectedPerSec** -- traces the selected rate of rejected bytes.

**TotalProduceRequestsPerSec** -- traces the selected rate of total produce requests.

**FailedProduceRequestsPerSec** -- traces the selected rate of failed produce requests.

**TotalFetchRequestsPerSec** -- traces the selected rate of total fetch requests.

**FailedFetchRequestsPerSec** -- traces the selected rate of failed fetch requests.

**Log Scale**

Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.

**Base at Zero**

Select to use zero (**0**) as the Y axis minimum for all graph traces.

**Time Range**

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .

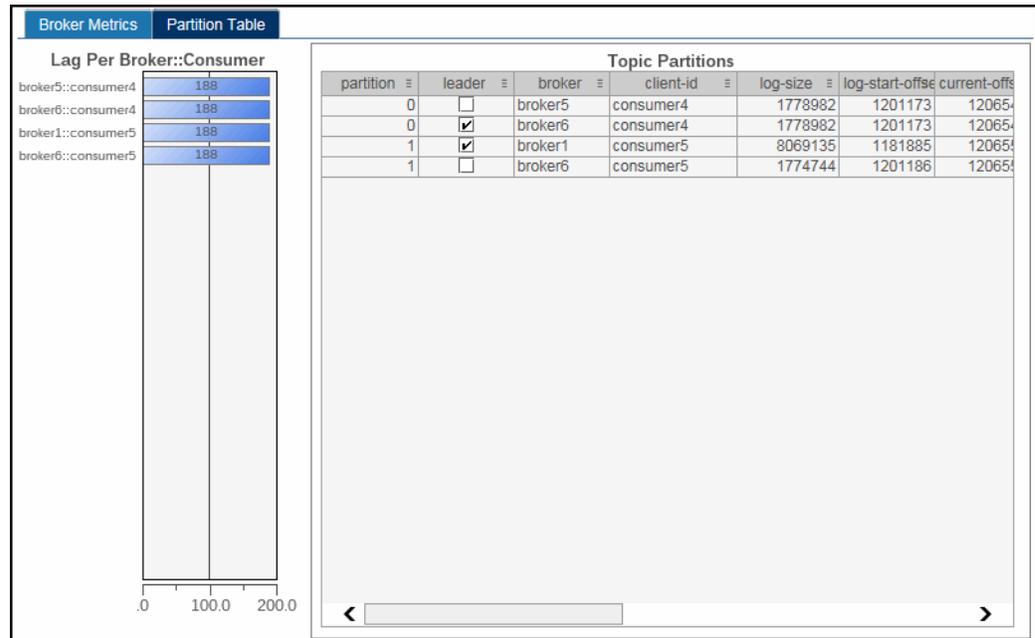


By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Partition Table Tab



### Lag Per Broker::Consumer

Lag per partition, where the partitions are identified by the broker hosting the partition and the consumer reading the partition.

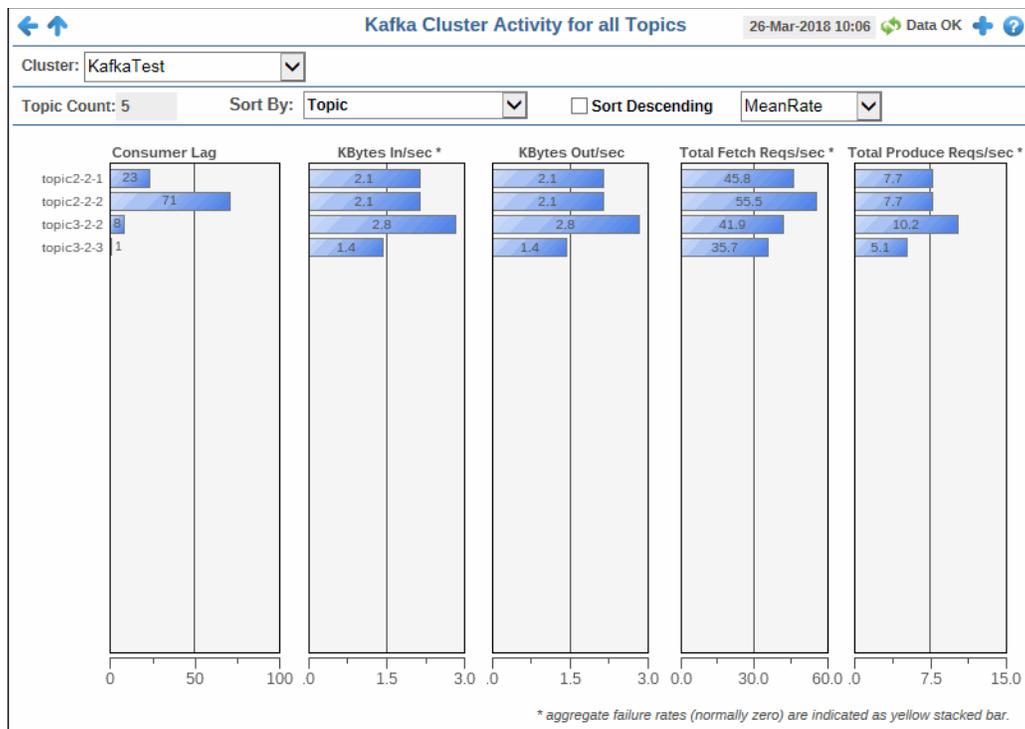
### Topic Partitions Table

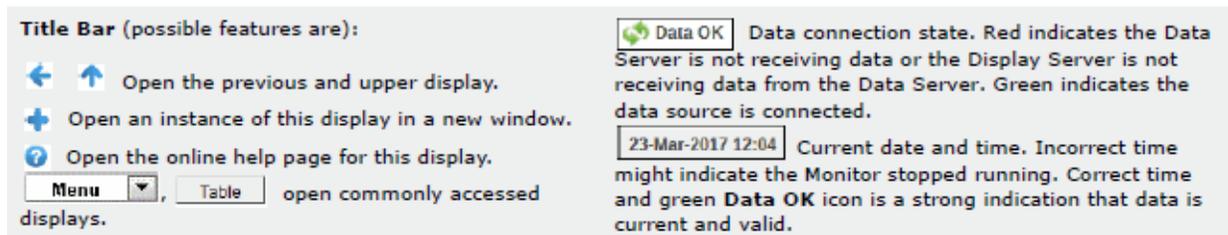
<b>cluster_name</b>	The name of the cluster.
<b>topic</b>	the name of the topic.
<b>Partition</b>	The name of the partition.
<b>Leader</b>	When checked, signifies that the broker is a leader on the partition.
<b>Broker</b>	The name of the broker.
<b>Client ID</b>	The ID of the consumer reading the topic.*
<b>broker_id</b>	The ID of the broker.
<b>Current Delta</b>	The difference between the current consumer position in the partition from the previous polling period to the current polling period.*
<b>Log Current Offset</b>	The offset of the message currently being consumed.*
<b>Current Rate</b>	The rate of change of the current consumer position.*
<b>Lag</b>	The difference between the current consumer position in the partition and the end of the log.*
<b>Lag Delta</b>	The difference in the amount of lag from the previous polling period to the current polling period.*
<b>Lag Rate</b>	The rate of change in the amount of lag.*

<b>Log End Delta</b>	The difference between the offset of the last message in the partition from the previous polling period to the current polling period.*
<b>Log End Offset</b>	The offset of the last message written to the log.*
<b>Log End Rate</b>	The rate of change of the last message offset.*
<b>Log Size</b>	The current number of messages in the log.*
<b>Log Start Offset</b>	The offset of the first message written to the log.
<b>time_stamp</b>	The date and time the row data was last updated.

## All Topics for Cluster

This display provides a view of the activity metrics on all topics for a particular cluster. You can view the metrics based on the mean rate, a 1 minute average rate, a 5 minute average rate, or a 15 minute average rate.





**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the selected cluster. Refer to KAFKA documentation for more information regarding these fields.

### Filter By:

- Cluster** Select the cluster for which you want to see data.
- <Rate Drop Down List>** Select the option for which you want to view data:
- Mean Rate** -- Select this option to view the average rate for each metric for the topics in the display.
  - 1 Minute Avg** -- Select this option to view the rate of incoming messages (per second) averaged over a one minute period for each metric for the topics in the display.
  - 5 Minute Avg** -- Select this option to view the rate of incoming messages (per second) averaged over a five minute period for each metric for the topics in the display.
  - 15 Minute Avg** -- Select this option to view the rate of incoming messages (per second) averaged over a fifteen minute period for each metric for the topics in the display.

**Topic Count** The number of topics found in the cluster.

**Sort By** Select the metric by which you want to sort the data in the display. When using this option with the **Sort Descending** toggle, the topics will be sorted in ascending or descending order using the option you select from this drop down. For example, if you select **Msgs In/sec** from this drop down and select the **Sort Descending** toggle, the topics listed in the display will be sorted so that the topic with the most **Msgs In/sec** will be listed at the top followed by the topic with the next most **Msgs In/sec**, and so on.

**Sort Descending** When toggled on, the topics listed in the display are sorted in descending order based on the selected metric in the **Sort By** drop down. When toggled off, the topics are listed in ascending order.

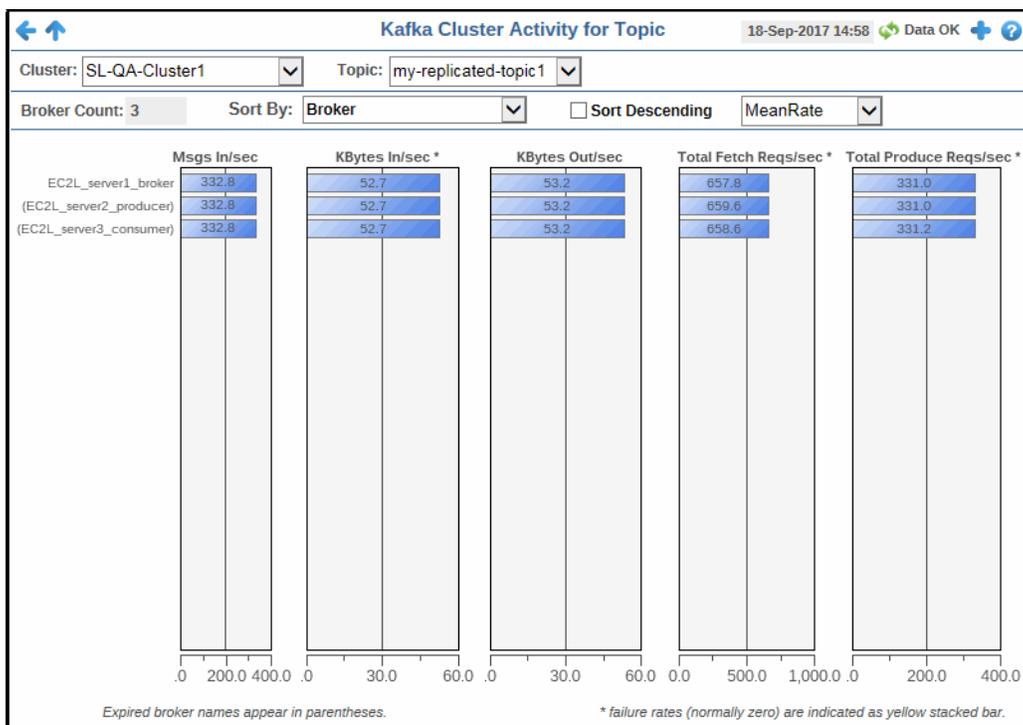
### Cluster Activity for Each Topic:

- Consumer Lag** The difference between the current consumer position in the partition and the end of the log.\*
- KBytes In/sec** The number of incoming kilobytes per second. For example, if you select **1 Minute Avg** from the drop down list, the average rate of incoming kilobytes per second for 1 minute.
- KBytes Out/sec** The number of outgoing kilobytes per second. For example, if you select **1 Minute Avg** from the drop down list, the average rate of outgoing kilobytes per second for 1 minute.

- Total Fetch Reqs/sec** The total number of fetch requests per second. For example, if you select **1 Minute Avg** from the drop down list, the average rate of fetch requests per second for 1 minute.
- Total Produce Reqs/sec** The total number of produce requests per second. For example, if you select **1 Minute Avg** from the drop down list, the average rate of producer requests per second for 1 minute.

### Single Topic for Cluster

This display provides a view of the activity metrics on all brokers for a particular topic. You can view the metrics based on the mean rate, a 1 minute average rate, a 5 minute average rate, or a 15 minute average rate.



**Title Bar (possible features are):**

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu Table open commonly accessed displays.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the selected cluster. Refer to KAFKA documentation for more information regarding these fields.

**Filter By:**

<b>Cluster</b>	Select the cluster for which you want to see data.
<b>Topic</b>	Select the topic for which you want to see data.
<b>&lt;Rate Drop Down List&gt;</b>	Select the option for which you want to view data: <b>Mean Rate</b> -- Select this option to view the average rate for each metric for the brokers in the display. <b>1 Minute Avg</b> -- Select this option to view the rate of incoming messages (per second) averaged over a one minute period for each metric for the brokers in the display. <b>5 Minute Avg</b> -- Select this option to view the rate of incoming messages (per second) averaged over a five minute period for each metric for the brokers in the display. <b>15 Minute Avg</b> -- Select this option to view the rate of incoming messages (per second) averaged over a fifteen minute period for each metric for the brokers in the display.

**Broker Count** The number of brokers found in the cluster with the associated topic.

**Sort By** Select the metric by which you want to sort the data in the display. When using this option with the **Sort Descending** toggle, the brokers will be sorted in ascending or descending order using the option you select from this drop down. For example, if you select **Msgs In/sec** from this drop down and select the **Sort Descending** toggle, the brokers listed in the display will be sorted so that the broker with the most **Msgs In/sec** will be listed at the top followed by the broker with the next most **Msgs In/sec**, and so on.

**Sort Descending** When toggled on, the brokers listed in the display are sorted in descending order based on the selected metric in the **Sort By** drop down. When toggled off, the brokers are listed in ascending order.

**Cluster Activity for Each Broker:**

<b>Msgs In/sec</b>	The number of incoming messages per second. For example, if you select <b>1 Minute Avg</b> from the drop down list, the average rate of incoming messages per second for 1 minute.
<b>KBytes In/sec</b>	The number of incoming kilobytes per second. For example, if you select <b>1 Minute Avg</b> from the drop down list, the average rate of incoming kilobytes per second for 1 minute.
<b>KBytes Out/sec</b>	The number of outgoing kilobytes per second. For example, if you select <b>1 Minute Avg</b> from the drop down list, the average rate of outgoing kilobytes per second for 1 minute.
<b>Total Fetch Reqs/sec</b>	The total number of fetch requests per second. For example, if you select <b>1 Minute Avg</b> from the drop down list, the average rate of fetch requests per second for 1 minute.
<b>Total Produce Reqs/sec</b>	The total number of produce requests per second. For example, if you select <b>1 Minute Avg</b> from the drop down list, the average rate of producer requests per second for 1 minute.

## Kafka Brokers View

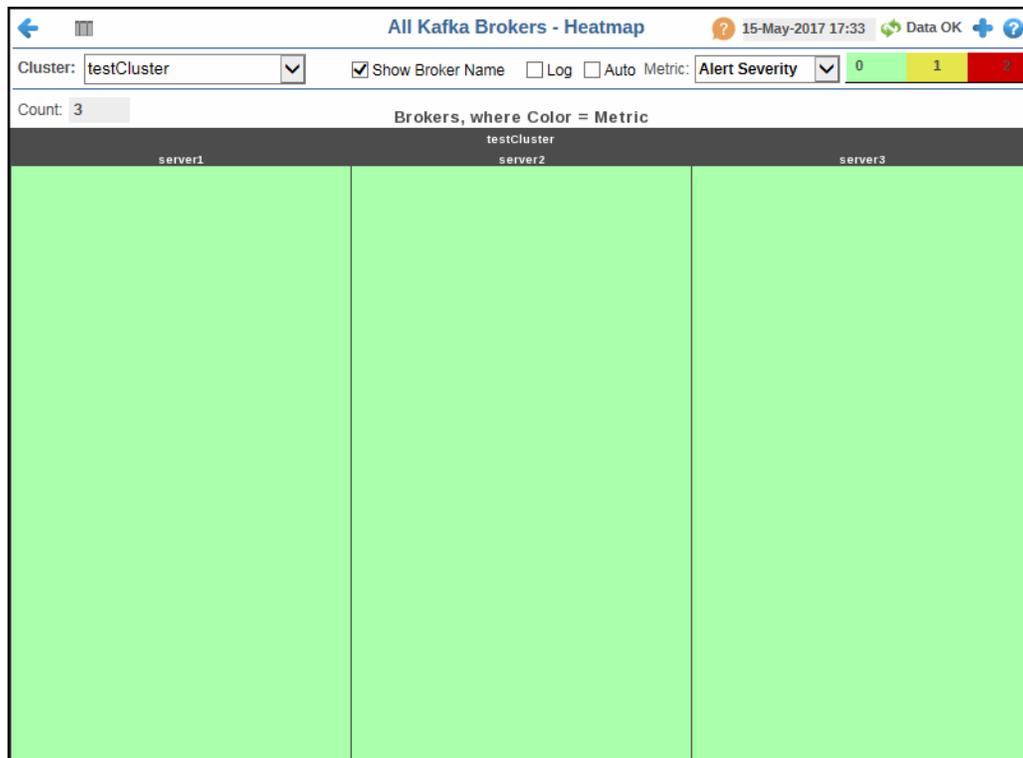
These displays provide detailed data for all brokers in heatmap and tabular form, provide details for all metrics for a particular broker in tabular form, and provide JVM runtime and broker status details for a particular broker. Displays in this View are:

- ["All Brokers Heatmap"](#): A heatmap view of all brokers in a heatmap format and their associated metrics.
- ["All Brokers Table"](#): A tabular view of your brokers and their associated metrics.
- ["All Broker Metrics Table"](#): A tabular and trend graph view of meter metrics, histogram metrics, and timer metrics for a particular broker.
- ["Single Broker Summary"](#): Contains JVM runtime data, broker status, topic, and topic trend details for a particular broker.

### All Brokers Heatmap

This heatmap display provides an easy-to-view interface that allows you to quickly identify the current status of each of your brokers for each available metric. You can view the brokers in the heatmap based on the following metrics: the current alert severity, the current alert count, the under replicated partitions count, the offline partitions count, the rate of incoming messages, the rate of incoming bytes, the rate of outgoing bytes, and the log flush latency value. By default, this display shows the heatmap based on the **Alert Severity** metric.

You can use the **Show Broker Name** check-box  to include or exclude labels in the heatmap, and you can mouse over a rectangle to see additional metrics for a broker. Clicking one of the rectangles in the heatmap opens the ["Single Broker Summary"](#) display, which allows you to see additional details for the selected broker.



**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

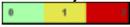
**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Fields and Data:**

- Cluster** Select the cluster for which you want to view data.
- Show Broker Name** Select this check box to display the names of the brokers at the top of each rectangle in the heatmap.
- Log** Select this check box to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Auto** Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value.  
**Note:** Some metrics auto-scale automatically, even when **Auto** is not selected.
- Metric** Choose a metric to view in the display.

**Alert Severity**

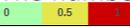
The current alert severity. Values range from **0** - **2**, as indicated in the color gradient  bar, where **2** is the highest Alert Severity:

-  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
-  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
-  Green indicates that no metrics have exceeded their alert thresholds.

**Alert Count**

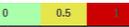
The total number of critical and warning unacknowledged alerts in the brokers. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.

**Under Replicated Partitions**

The number of under-replicated partitions. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **KafkaBrokerUnderReplicatedPartns**. The middle value in the gradient bar indicates the middle value of the range.

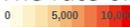
When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

**Offline Partitions**

The number of offline partitions. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **KafkaBrokerOfflinePartitionCnt**. The middle value in the gradient bar indicates the middle value of the range.

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

**Msgs In Per Sec**

The rate of incoming messages (per second). The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **KafkaBrokerMsgsInPerSec**. The middle value in the gradient bar indicates the middle value of the range.

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

**Bytes In Per Sec**

The rate of incoming bytes (per second). The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **KafkaBrokerBytesInPerSec**. The middle value in the gradient bar indicates the middle value of the range.

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

**Bytes Out Per Sec**

The rate of outgoing bytes (per second). The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **KafkaBrokerBytesOutPerSec**. The middle value in the gradient bar indicates the middle value of the range.

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

**Log Flush Latency 95 Pctile**

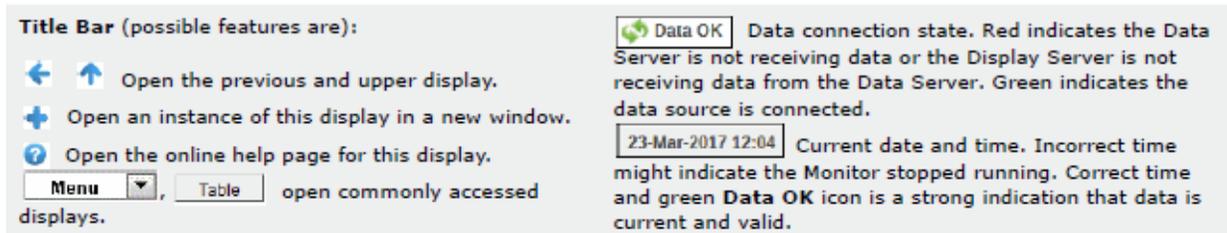
The log flush latency for the top five percent of values. The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **KafkaBrokerLogFlushLatency95P**. The middle value in the gradient bar indicates the middle value of the range.

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

## All Brokers Table

The table in this display provides a view of all of your brokers and their associated metric data including cluster name, broker name, broker ID, alert level, alert count, and the current value of each gathered metric. You can click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by clicking a row to view details for the selected broker in the ["Single Broker Summary"](#) display.

All Kafka Brokers - Table							
Cluster: All Clusters							
Count: 3							
Kafka Brokers							
Cluster Name	Broker Name	Broker ID	Alert Level	Alert Count	Broker State	Active Controller	CPU %
KafkaTest	broker1	0	0	0	RunningAsBroker	<input checked="" type="checkbox"/>	12.02
KafkaTest	broker2	1	0	0	RunningAsBroker	<input type="checkbox"/>	1.55
KafkaTest	broker3	2	0	0	RunningAsBroker	<input type="checkbox"/>	5.96



**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the selected cluster. Refer to KAFKA documentation for more information regarding these fields.

### Filter By

**Cluster** Select the cluster for which you want to view data.

**Count** Lists the number of brokers found as a result of the cluster that was selected and displayed in the Kafka Brokers table.

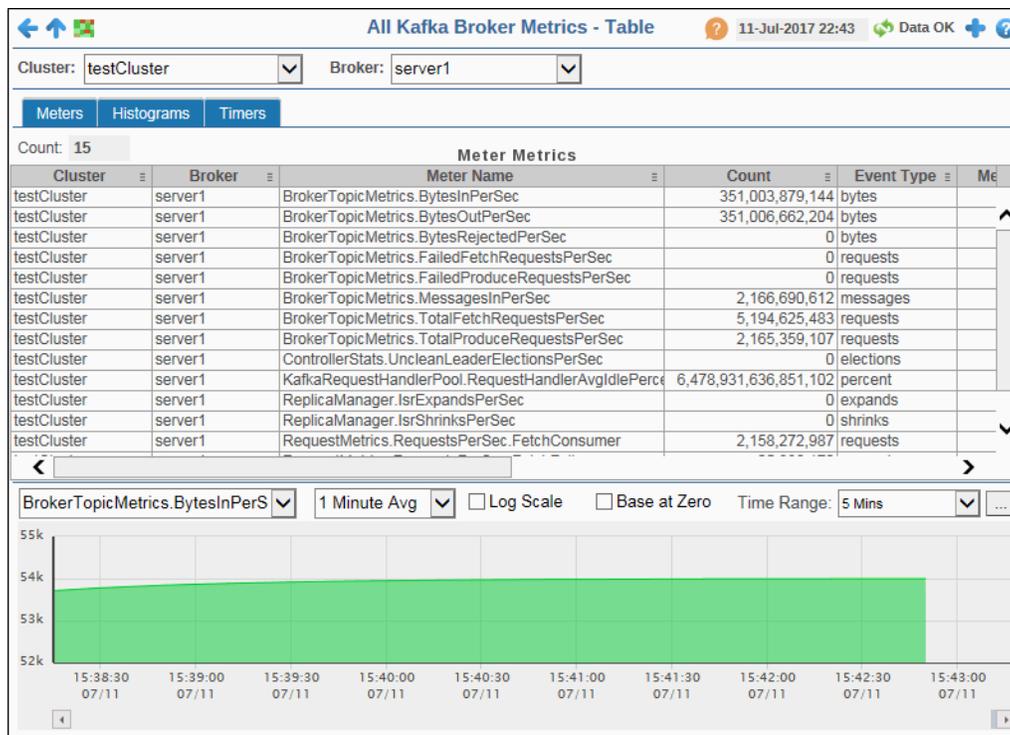
### Kafka Brokers Table:

<b>Cluster Name</b>	The name of the cluster.
<b>Broker Name</b>	The name of the broker.
<b>Broker ID</b>	The broker ID for the server.
<b>Alert Level</b>	The current alert severity.  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
<b>Alert Count</b>	The total number of alerts for the host.
<b>Broker State</b>	The current state of the kafka broker.*
<b>Active Controller</b>	Denotes whether the broker is an active controller.*
<b>CPU %</b>	The percentage of CPU being used by the broker.*
<b>Mem Used %</b>	The percentage of JVM memory being used by the broker.*
<b>Leader Count</b>	The number of leaders on the broker.*
<b># Partitions</b>	The number of partitions on the broker.*
<b># Offline Partitions</b>	The number of partitions without an active leader on the broker.*
<b>Under Replicated Partitions</b>	The number of partition replicas that are out of sync (total number of replicas minus the total number of in-sync replicas) on the broker.*

<b>Preferred Replica Imbalance Count</b>	The number of topics whose replicas are not balanced on the broker.*
<b>ZK Disconnect Rate</b>	The mean rate of zookeeper disconnects per broker, in seconds.
<b># Purgatory Fetch</b>	The number of fetch requests currently in purgatory (and waiting to be satisfied).*
<b># Purgatory Heartbeat</b>	The number of requests in purgatory due to failed heartbeat tests.*
<b># Purgatory Produce</b>	The number of produce requests currently in purgatory (and waiting to be satisfied).*
<b># Purgatory Rebalance</b>	The number of changes that need to be propagated to the replicas so that the partitions are no longer in purgatory.*
<b># Purgatory Topic</b>	The number of requests (based on topics) currently in purgatory.*
<b>Network Processor Avg % Idle</b>	The average fraction of time the network processors are idle.*
<b>Version</b>	The current version of Kafka.*
<b>JMX Connection String</b>	The JMX connection string used.*
<b>Connected?</b>	Denotes whether or not the broker is connected.*
<b>Expired</b>	<p>When checked, performance data in the row has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the RTView Configuration Application &gt; (<b>KAFKAMON-LOCAL/Project Name</b>) &gt; <b>Solution Package Configuration</b> &gt; <b>Apache Kafka</b> &gt; <b>DATA STORAGE</b> &gt; <b>Duration</b> &gt; <b>Expire Time</b> property. The RTView Configuration Application &gt; (<b>KAFKAMON-LOCAL/Project Name</b>) &gt; <b>Solution Package Configuration</b> &gt; <b>Apache Kafka</b> &gt; <b>DATA Storage</b> &gt; <b>Duration</b> &gt; <b>Delete Time</b> property allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.</p> <p>For example, if <b>Expire Time</b> was set to 120 and <b>Delete Time</b> was set to 3600, then the <b>Expired</b> check box would be checked after 120 seconds and the row would be removed from the table after 3600 seconds.</p>
<b>Timestamp</b>	The date and time the row data was last updated.

## All Broker Metrics Table

This display contains broker metrics broken down into three tabs: **Meters**, **Histograms**, and **Timers**. The table in the **Meters** tab provides a view of all of your metered metrics by broker and their associated metric data including count, event type, and rate data. The table in the **Histograms** tab provides a view of the histogram metrics for the selected broker. The table in the **Timers** tab provides a view of all the timers for the selected broker and their associated metrics. Each of the tabs also contains a trend graph, which provides a trend chart for each of the metrics listed in the associated table.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

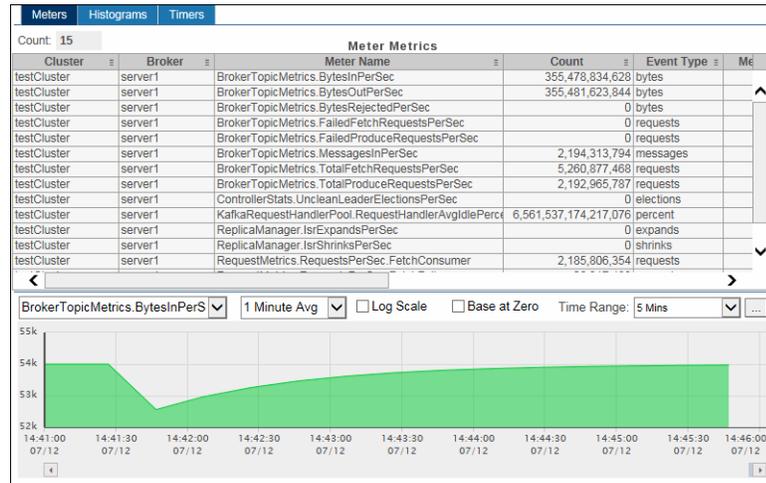
Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the selected cluster. Refer to KAFKA documentation for more information regarding these fields.

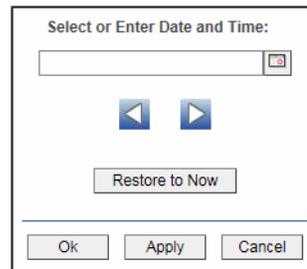
**Filter By:**

- Cluster** Select the cluster containing the broker for which you want to see data.
- Broker** Select the broker for which you want to see data.
- Count** The number of meters/histograms/timers found using the filter.

**Meters Tab:**

- Cluster** The name of the cluster.
- Broker** The name of the broker.
- Meter Name** The name of the metered metric.
- Count** The total count for the metered metric.
- Event Type** The type (unit) of metered metric.\*
- Mean Rate** The average rate for the meter, based on the **Rate Unit**.\*
- 1 Minute Avg** The rate of incoming messages (per second) averaged over a one minute period for the meter, based on the **Rate Unit**.\*
- 5 Minute Avg** The rate of incoming messages (per second) averaged over a five minute period for the meter, based on the **Rate Unit**.\*
- 15 Minute Avg** The rate of incoming messages (per second) averaged over a fifteen minute period for the meter, based on the **Rate Unit**.\*
- Rate Unit** The unit of measure used to calculate **1 Minute Avg**, **5 Minute Avg**, **15 Minute Avg**, and **Mean Rate**.\*
- Expired** When checked, performance data in the row has not been received within the time specified (in seconds) in the **Expire Time** field in the RTView Configuration Application > (**KAFKAMON-LOCAL/Project Name**) > **Solution Package Configuration** > **Apache Kafka** > **DATA STORAGE** > **Duration** > **Expire Time** property. The RTView Configuration Application > (**KAFKAMON-LOCAL/Project Name**) > **Solution Package Configuration** > **Apache Kafka** > **DATA Storage** > **Duration** > **Delete Time** property allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
- For example, if **Expire Time** was set to 120 and **Delete Time** was set to 3600, then the **Expired** check box would be checked after 120 seconds and the row would be removed from the table after 3600 seconds.

<b>Timestamp</b>	The date and time the row data was last updated.
<b>Metric Performance Trends</b>	The trend chart provides a moving chart over the selected time range for each of the <b>Meter Names</b> listed in the <b>Meters Metrics</b> table.
<b>&lt;Meter Names Drop Down List&gt;</b>	Select the meter name for which you want to view data.
<b>&lt;Rate Drop Down List&gt;</b>	Select the option for which you want to view data: <b>Mean Rate</b> -- Select this option to view the average rate for each metric for the metrics in the display. <b>1 Minute Avg</b> -- Select this option to view the rate of incoming messages (per second) averaged over a one minute period for each metric for the metrics in the display. <b>5 Minute Avg</b> -- Select this option to view the rate of incoming messages (per second) averaged over a five minute period for each metric for the metrics in the display. <b>15 Minute Avg</b> -- Select this option to view the rate of incoming messages (per second) averaged over a one minute period for each metric for the metrics in the display.
<b>Log Scale</b>	Select to enable a logarithmic scale. Use <b>Log Scale</b> to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. <b>Log Scale</b> makes data on both scales visible by applying logarithmic values rather than actual values to the data.
<b>Base at Zero</b>	Select to use zero (0) as the Y axis minimum for all graph traces.
<b>Time Range</b>	Select a time range from the drop down menu varying from <b>2 Minutes</b> to <b>Last 7 Days</b> , or display <b>All Data</b> . To specify a time range, click Calendar  .

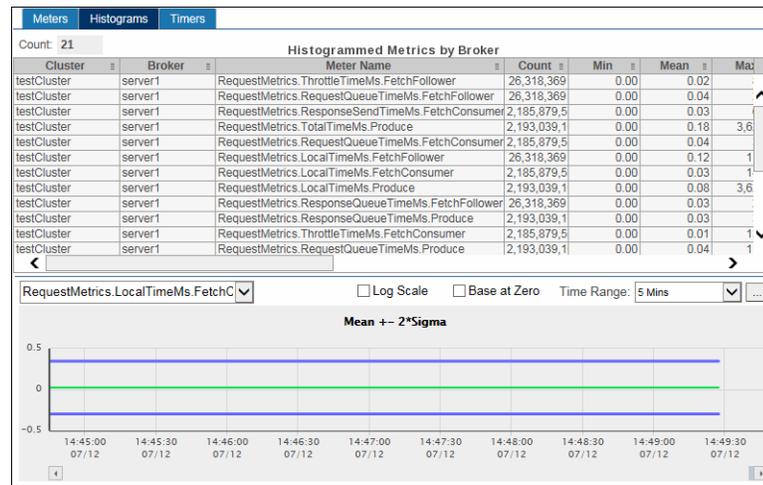


By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

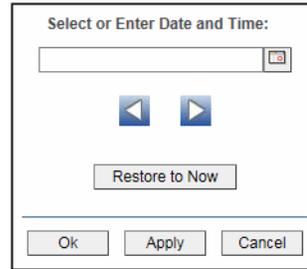
## Histograms Tab



<b>Cluster</b>	The name of the cluster.
<b>Broker</b>	The name of the broker.
<b>Meter Name</b>	The name of the metered metric.
<b>Count</b>	The total count for the metered metric.*
<b>Min</b>	The minimum number of occurrences for the meter during the current polling period.*
<b>Mean</b>	The average number of occurrences for the meter during the current polling period.*
<b>Max</b>	The maximum number of occurrences for the meter during the current polling period.*
<b>Std Dev</b>	The standard deviation for the number of occurrences for the meter during the current polling period.*
<b>50th Percentile</b>	The 50th percentile value for the number of occurrences for the meter during the current polling period.*
<b>75th Percentile</b>	The 75th percentile value for the number of occurrences for the meter during the current polling period.*
<b>95th Percentile</b>	The 95th percentile value for the number of occurrences for the meter during the current polling period.*
<b>98th Percentile</b>	The 98th percentile value for the number of occurrences for the meter during the current polling period.*
<b>99th Percentile</b>	The 99th percentile value for the number of occurrences for the meter during the current polling period.*
<b>999th Percentile</b>	The 999th percentile value for the number of occurrences for the meter during the current polling period.*

<b>Expired</b>	<p>When checked, performance data in the row has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the RTView Configuration Application &gt; (<b>KAFKAMON-LOCAL/Project Name</b>) &gt; <b>Solution Package Configuration</b> &gt; <b>Apache Kafka</b> &gt; <b>DATA STORAGE</b> &gt; <b>Duration</b> &gt; <b>Expire Time</b> property. The RTView Configuration Application &gt; (<b>KAFKAMON-LOCAL/Project Name</b>) &gt; <b>Solution Package Configuration</b> &gt; <b>Apache Kafka</b> &gt; <b>DATA Storage</b> &gt; <b>Duration</b> &gt; <b>Delete Time</b> property allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.</p> <p>For example, if <b>Expire Time</b> was set to 120 and <b>Delete Time</b> was set to 3600, then the <b>Expired</b> check box would be checked after 120 seconds and the row would be removed from the table after 3600 seconds.</p>
<b>Timestamp</b>	The date and time the row data was last updated.
<b>Metric Performance Trends</b>	<p>The trend chart provides a moving chart over the selected time range for each of the <b>Meter Names</b> listed in the <b>Histogrammed Metrics by Broker</b> table.</p> <p><b>&lt;Meter Names Drop Down List&gt;</b> Select the meter name for which you want to view data.</p> <p><b>&lt;Rate Drop Down List&gt;</b> Select the option for which you want to view data:</p> <p><b>Mean Rate</b> -- Select this option to view the average rate for each metric for the metrics in the display.</p> <p><b>1 Minute Avg</b> -- Select this option to view the rate of incoming messages (per second) averaged over a one minute period for each metric for the metrics in the display.</p> <p><b>5 Minute Avg</b> -- Select this option to view the rate of incoming messages (per second) averaged over a five minute period for each metric for the metrics in the display.</p> <p><b>15 Minute Avg</b> -- Select this option to view the rate of incoming messages (per second) averaged over a fifteen minute period for each metric for the metrics in the display.</p> <p><b>Log Scale</b> Select to enable a logarithmic scale. Use <b>Log Scale</b> to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. <b>Log Scale</b> makes data on both scales visible by applying logarithmic values rather than actual values to the data.</p> <p><b>Base at Zero</b> Select to use zero (<b>0</b>) as the Y axis minimum for all graph traces.</p>

**Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar  .

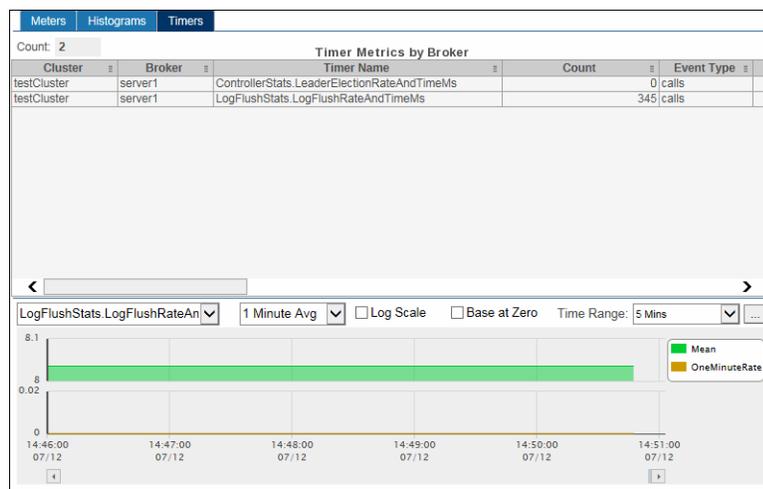


By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

**Timers Tab:**



- Cluster** The name of the cluster.
- Broker** The name of the broker.
- Timer Name** The name of the timer.
- Count** The total count for the timer.
- Event Type** The event type for the timer.
- Mean** The average number of events for the timer during the current polling period.\*
- Mean Rate** The average rate (based on **Rate Unit**) of events for the timer during the current polling period.\*
- Min** The minimum number of events for the timer during the current polling period.\*
- Max** The maximum number of events for the timer during the current polling period.\*

<b>Std Dev</b>	The standard deviation for the number of events for the timer during the current polling period.*
<b>1 Minute Avg</b>	The rate of incoming messages (per second) averaged over a one minute period for the timer, based on the <b>Rate Unit</b> .*
<b>5 Minute Avg</b>	The rate of incoming messages (per second) averaged over a five minute period for the timer, based on the <b>Rate Unit</b> .*
<b>15 Minute Avg</b>	The rate of incoming messages (per second) averaged over a fifteen minute period for the timer, based on the <b>Rate Unit</b> .*
<b>50th Percentile</b>	The 50th percentile value for the number of events for the timer during the current polling period.*
<b>75th Percentile</b>	The 75th percentile value for the number of events for the timer during the current polling period.*
<b>95th Percentile</b>	The 95th percentile value for the number of events for the timer during the current polling period.*
<b>98th Percentile</b>	The 98th percentile value for the number of events for the timer during the current polling period.*
<b>999th Percentile</b>	The 999th percentile value for the number of events for the timer during the current polling period.*
<b>99th Percentile</b>	The 99th percentile value for the number of events for the timer during the current polling period.*
<b>Latency Unit</b>	The unit of measure used to calculate the latency.*
<b>Rate Unit</b>	The unit of measure used to calculate <b>1 Minute Rate</b> , <b>5 Minute Rate</b> , <b>15 Minute Rate</b> , and <b>Mean Rate</b> .*
<b>Expired</b>	<p>When checked, performance data in the row has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the RTView Configuration Application &gt; (<b>KAFKAMON-LOCAL/Project Name</b>) &gt; <b>Solution Package Configuration</b> &gt; <b>Apache Kafka</b> &gt; <b>DATA STORAGE</b> &gt; <b>Duration</b> &gt; <b>Expire Time</b> property. The RTView Configuration Application &gt; (<b>KAFKAMON-LOCAL/Project Name</b>) &gt; <b>Solution Package Configuration</b> &gt; <b>Apache Kafka</b> &gt; <b>DATA Storage</b> &gt; <b>Duration</b> &gt; <b>Delete Time</b> property allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.</p> <p>For example, if <b>Expire Time</b> was set to 120 and <b>Delete Time</b> was set to 3600, then the <b>Expired</b> check box would be checked after 120 seconds and the row would be removed from the table after 3600 seconds.</p>
<b>Timestamp</b>	The date and time the row data was last updated.
<b>Timers Performance Trends</b>	<p>The trend chart provides a moving chart over the selected time range for each of the <b>Timer Names</b> listed in the <b>Timer Metrics by Broker</b> table.</p> <p><b>&lt;Timer Names Drop Down List&gt;</b> Select the timer name for which you want to view data.</p>

**<Rate Drop  
Down List>**

Select the option for which you want to view data:

**Mean Rate** -- Select this option to view the average rate for each metric for the metrics in the display.

**1 Minute Avg** -- Select this option to view the rate of incoming messages (per second) averaged over a one minute period for each metric for the metrics in the display.

**5 Minute Avg** -- Select this option to view the rate of incoming messages (per second) averaged over a five minute period for each metric for the metrics in the display.

**15 Minute Avg** -- Select this option to view the rate of incoming messages (per second) averaged over a fifteen minute period for each metric for the metrics in the display.

**Log Scale**

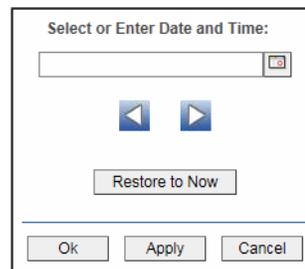
Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.

**Base at Zero**

Select to use zero (**0**) as the Y axis minimum for all graph traces.

**Time Range**

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



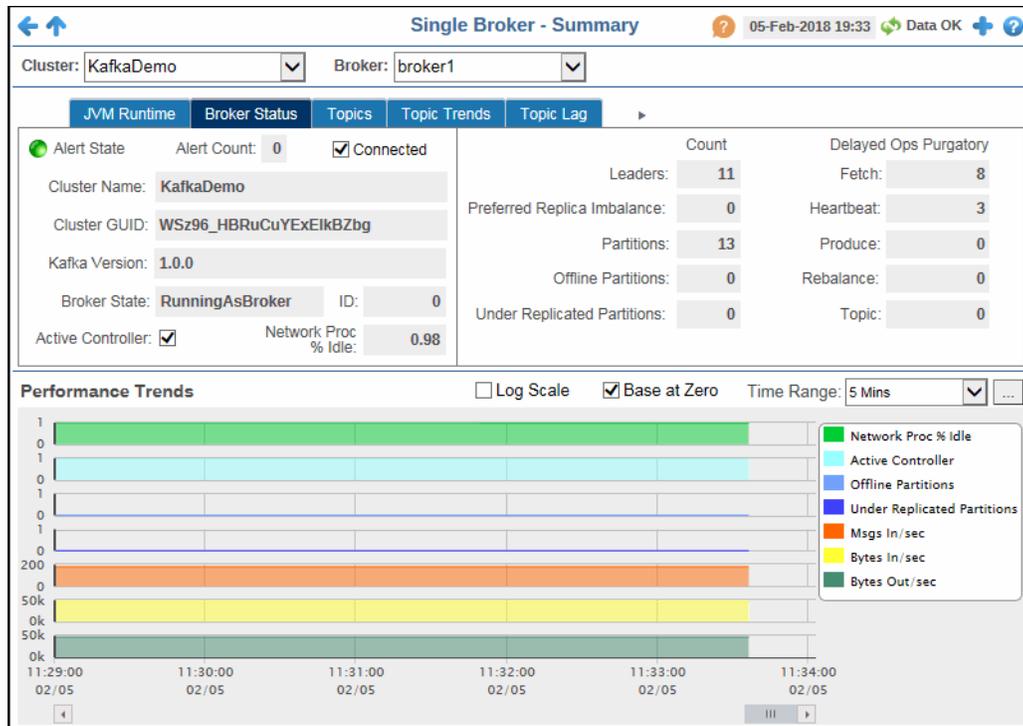
By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Single Broker Summary

This display provides a view of the current and historical metrics for a single broker, including JVM runtime data, broker status, topic data, and topic trend data.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the selected cluster. Refer to KAFKA documentation for more information regarding these fields.

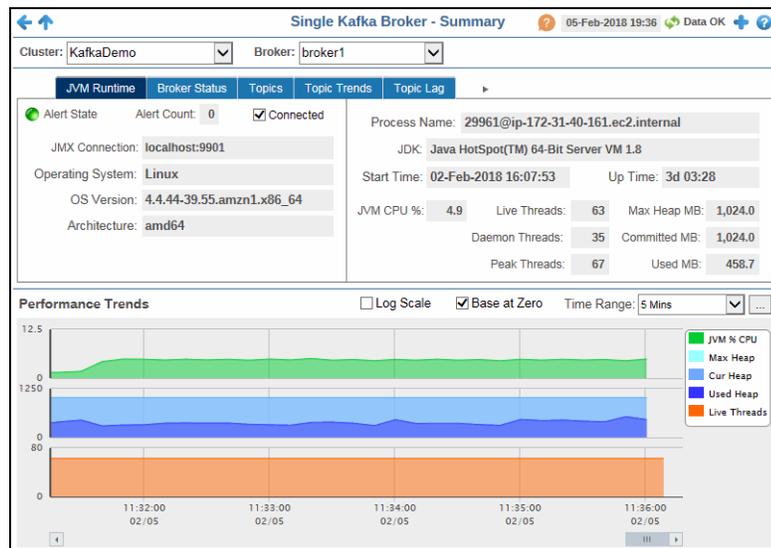
### Filter By:

The display might include these filtering options:

- Cluster** Select the cluster for which you want to show data in the display.
- Broker** Select the broker for which you want to show data in the display.

**Topic**

Only displays when the **Topic Trends** tab is selected. Select the topic for which you want to show data in the display.

**JVM Runtime Tab****Alert State**

The current alert severity.

- Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
- Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
- Green indicates that no metrics have exceeded their alert thresholds.

**Alert Count**

The total number of current alerts.

**Connected**

When checked, denotes that the broker is connected.

**JMX Connection**

The name of the JMX connection.\*

**Operating System**

The operating system installed on the broker.\*

**OS Version**

The version number of the operating system.\*

**Architecture**

The type of processor being used.\*

**Process Name**

The name of the process.\*

**JDK**

The JDK version number.\*

**Start Time**

The date and time when the broker was started.\*

**Up Time**

The amount of time the broker has been up and running.\*

**JVM CPU %**

The percentage of CPU used by the JVM.\*

**Live Threads**

The number of live threads on the broker.\*

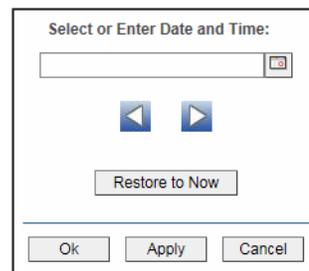
**Max Heap MB**

The maximum amount of available heap, in megabytes.\*

**Daemon Threads**

The number of daemon threads running on the broker.\*

- Committed MB** The total number of megabytes committed on the broker.\*
- Peak Threads** The highest number of threads running at one time during the current polling period.\*
- Used MB** The number of used megabytes on the broker.\*
- Performance Trends** Traces the following:
- JVM % CPU** -- traces the percentage of CPU used by the JVM.
  - Max Heap** -- traces the maximum amount of available heap.
  - Cur Heap** -- traces the current amount of heap being used.
  - Used Heap** -- traces the highest amount of heap used.
  - Live Threads** -- traces the number of live threads.
- Log Scale** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Base at Zero** Select to use zero (0) as the Y axis minimum for all graph traces.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .

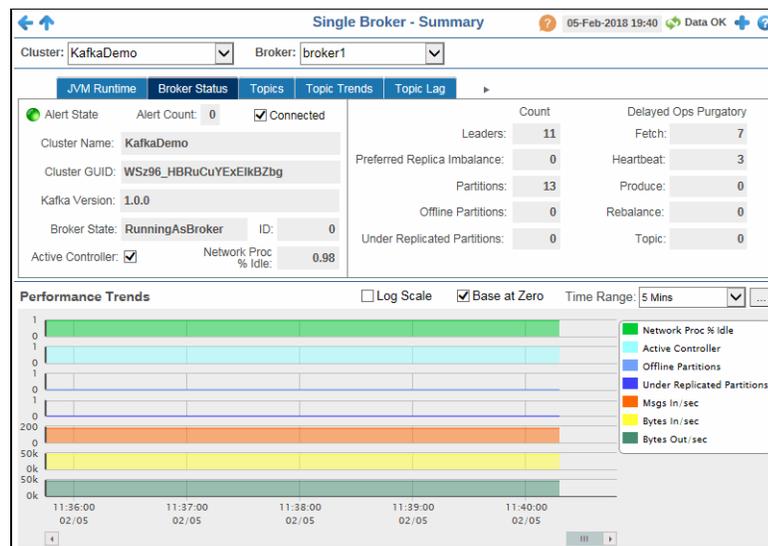


By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Broker Status Tab



### Alert State

The current alert severity.

- Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
- Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
- Green indicates that no metrics have exceeded their alert thresholds.

### Alert Count

The total number of current alerts.

### Connected

When checked, denotes that the broker is connected.

### Cluster Name

The name of the cluster in which the broker is contained.

### Cluster GUID

Lists the cluster's globally unique identifier.

**Note:** This field will not be populated for brokers running on Kafka Version 0.9.\*, and the **KafkaClusterSplitBrain** alert will not work properly for those brokers.

### Kafka Version

The current version of Kafka installed on the broker.

### Broker State

The current state of the Kafka broker.

### ID

The broker ID for the server.

### Active Controller

Denotes whether or not the broker is an active controller.

### Network Proc % Idle

The average fraction of time the network processors are idle.\*

### Count

#### Leaders

The number of leaders on the broker.\*

#### Preferred Replica Imbalance

The number of topics whose replicas are not balanced on the broker.\*

#### Partitions

The number of partitions on the broker.\*

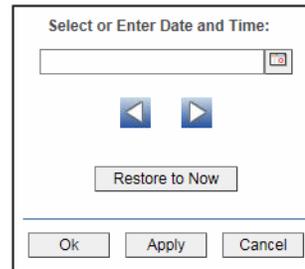
#### Offline Partitions

The number of partitions on the broker that are currently offline.\*

<b>Under Replicated Partitions</b>	The number of partition replicas that are not in sync on the broker.*
<b>Delayed Ops Purgatory</b>	
<b>Fetch</b>	The number of fetch requests currently in purgatory (and waiting to be satisfied).*
<b>Heartbeat</b>	The number of requests in purgatory due to failed heartbeat tests.*
<b>Produce</b>	The number of produce requests currently in purgatory (and waiting to be satisfied).*
<b>Rebalance</b>	The frequency with which the partition rebalance check is triggered by the controller.*
<b>Topic</b>	The number of requests (based on topics) currently in purgatory.*
<b>Performance Trends</b>	Traces the following: <ul style="list-style-type: none"> <li><b>Network Proc % Idle</b> -- traces the average fraction of time the network processors are idle.</li> <li><b>Active Controller</b> -- traces whether or not the broker is/was an active controller.</li> <li><b>Offline Partitions</b> -- traces the number of offline partitions.</li> <li><b>Under Replicated Partitions</b> -- traces the number of partition replicas out of sync on the broker.</li> <li><b>Msgs In/sec</b> -- traces the number of incoming messages per second.</li> <li><b>Bytes In/sec</b> -- traces the number of incoming bytes per second.</li> <li><b>Bytes Out/sec</b> -- traces the number of outgoing bytes per second.</li> </ul>
<b>Log Scale</b>	Select to enable a logarithmic scale. Use <b>Log Scale</b> to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. <b>Log Scale</b> makes data on both scales visible by applying logarithmic values rather than actual values to the data.
<b>Base at Zero</b>	Select to use zero ( <b>0</b> ) as the Y axis minimum for all graph traces.

**Time Range**

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .

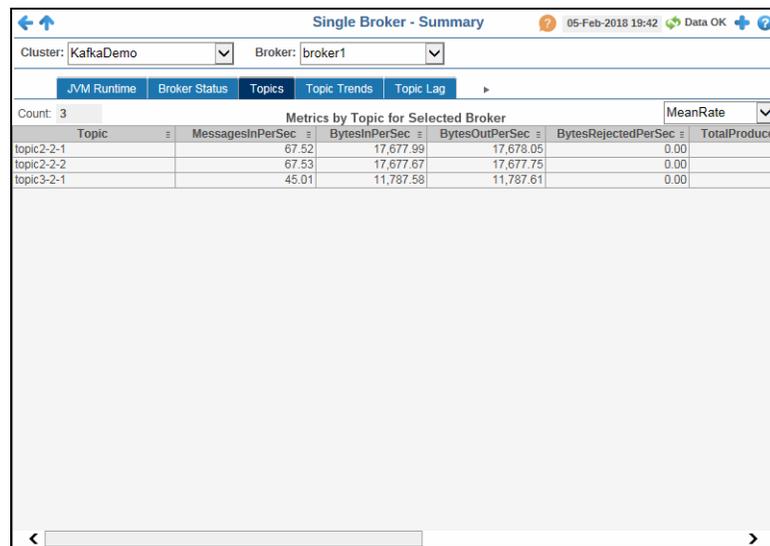


The dialog box titled "Select or Enter Date and Time:" contains a text input field with a calendar icon on the right. Below the input field are two navigation arrows (left and right). At the bottom of the dialog are three buttons: "Restore to Now", "Ok", "Apply", and "Cancel".

By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

**Topics Tab**


The screenshot shows the "Single Broker - Summary" page for cluster "KafkaDemo" and broker "broker1". The "Topics" tab is selected, showing a table of metrics for 3 topics. The table has columns for Topic, MessagesInPerSec, BytesInPerSec, BytesOutPerSec, BytesRejectedPerSec, and TotalProduce. A "MeanRate" dropdown is visible on the right side of the table.

Topic	MessagesInPerSec	BytesInPerSec	BytesOutPerSec	BytesRejectedPerSec	TotalProduce
topic2-2-1	67.52	17,677.99	17,678.05	0.00	
topic2-2-2	67.53	17,677.67	17,677.75	0.00	
topic3-2-1	45.01	11,787.58	11,787.61	0.00	

**Count**

The total number of topics listed in the table.

**<Rate Drop Down List>**

Select the option for which you want to view data.

**Mean Rate**

Select this option to view the average rate for each metric for the topics in the display.

**One Minute**

Select this option to view the 1 minute rate for each metric for the topics in the display.

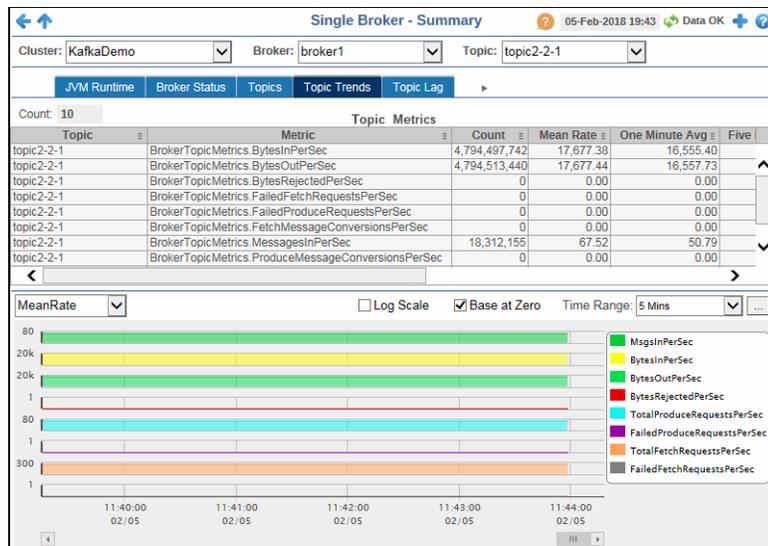
**Five Minute**

Select this option to view the 5 minute rate for each metric for the topics in the display.

**Fifteen Minute** Select this option to view the 15 minute rate for each metric for the topics in the display.

- Topic** Lists the name of the topic.
- Messages In Per Sec** The rate of incoming messages
- Bytes In Per Sec** The rate of incoming bytes
- Bytes Out Per Sec** The rate of outgoing bytes.
- Bytes Rejected Per Sec** The rate of rejected bytes.
- Total Produce Requests Per Sec** The rate of total produce requests.
- Failed Produce Requests Per Sec** The rate of failed produce requests.
- Total Fetch Requests Per Sec** The rate of total fetch requests.
- Failed Fetch Requests Per Sec** The rate of failed fetch requests.

**Topic Trends Tab**

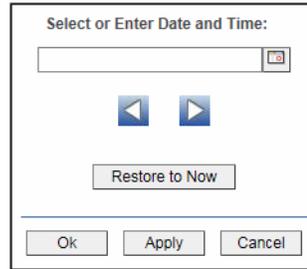


- Count** The total number of topic metrics listed in the table.
- Topic** The name of the topic
- Metric** The name of the metric.
- Count** The total number of the particular metric.
- Mean Rate** The mean rate of the metric.

<b>One Minute Avg</b>	The (one minute) rate for the metric, based on the Rate Units.
<b>Five Minute Avg</b>	The (five minute) rate for the metric, based on the Rate Units.
<b>Fifteen Minute Avg</b>	The (fifteen minute) rate for the metric, based on the Rate Units.
<b>Event Type</b>	The event type for the topic metric.
<b>Rate Units</b>	The unit of measure used to calculate the <b>Mean Rate</b> , the <b>One Minute Avg</b> , the <b>Five Minute Avg</b> , and the <b>Fifteen Minute Avg</b> .
<b>Expired</b>	<p>When checked, performance data in the row has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the RTView Configuration Application &gt; (<b>KAFKAMON-LOCAL/Project Name</b>) &gt; <b>Solution Package Configuration</b> &gt; <b>Apache Kafka</b> &gt; <b>DATA STORAGE</b> &gt; <b>Duration</b> &gt; <b>Expire Time</b> property. The RTView Configuration Application &gt; (<b>KAFKAMON-LOCAL/Project Name</b>) &gt; <b>Solution Package Configuration</b> &gt; <b>Apache Kafka</b> &gt; <b>DATA Storage</b> &gt; <b>Duration</b> &gt; <b>Delete Time</b> property allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.</p> <p>For example, if <b>Expire Time</b> was set to 120 and <b>Delete Time</b> was set to 3600, then the <b>Expired</b> check box would be checked after 120 seconds and the row would be removed from the table after 3600 seconds.</p>
<b>time_stamp</b>	The date and time the row data was last updated.
<b>Rate Trends</b>	<p>Select the option for which you want to view data:</p> <p><b>MeanRate</b> -- Select this option to view the average rate for each metric in the trend graph.</p> <p><b>One Minute</b>-- Select this option to view the rate averaged over a one minute period for each metric in the trend graph.</p> <p><b>Five Minute</b> -- Select this option to view the rate averaged over a five minute period for each metric in the trend graph.</p> <p><b>Fifteen Minute</b> -- Select this option to view the rate averaged over a fifteen minute period for each metric in the trend graph.</p> <p>Traces the following:</p> <ul style="list-style-type: none"> <li><b>MsgsInPerSec</b> -- traces the selected rate of incoming messages.</li> <li><b>BytesInPerSec</b> -- traces the selected rate of incoming bytes.</li> <li><b>BytesOutPerSec</b> -- traces the selected rate of outgoing bytes.</li> <li><b>BytesRejectedPerSec</b> -- traces the selected rate of rejected bytes.</li> <li><b>TotalProduceRequestsPerSec</b> -- traces the selected rate of total produce requests.</li> <li><b>FailedProduceRequestsPerSec</b> -- traces the selected rate of failed produce requests.</li> <li><b>TotalFetchRequestsPerSec</b> -- traces the selected rate of total fetch requests.</li> <li><b>FailedFetchRequestsPerSec</b> -- traces the selected rate of failed fetch requests.</li> </ul> <p><b>Log Scale</b>      Select to enable a logarithmic scale. Use <b>Log Scale</b> to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. <b>Log Scale</b> makes data on both scales visible by applying logarithmic values rather than actual values to the data.</p> <p><b>Base at Zero</b>    Select to use zero (<b>0</b>) as the Y axis minimum for all graph traces.</p>

**Time Range**

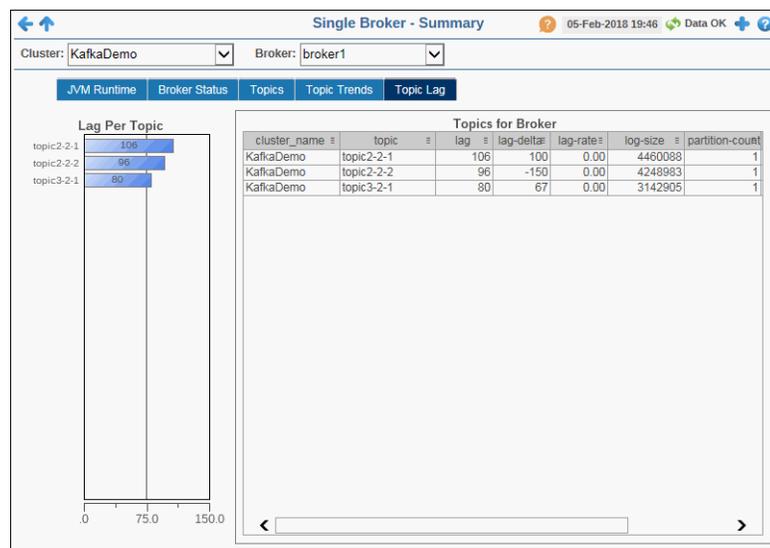
Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

**Topic Lag Tab****Lag Per Topic Bar Graph**

Displays the lag per topic in a bar graph format.

**Topics for Broker Table**

**cluster\_name** The name of the cluster in which the topic resides.

**topic** The name of the topic.

**lag** The difference between the current consumer position in the partition and the end of the log.\*

**lag-delta** The difference in the amount of lag from the previous polling period to the current polling period.\*

<b>lag-rate</b>	The rate of change in the amount of lag.*
<b>log-size</b>	The current number of messages in the log.*
<b>partition-count</b>	The number of partitions containing the topic.
<b>time_stamp</b>	The date and time the row data was last updated.

## Kafka Zookeepers View

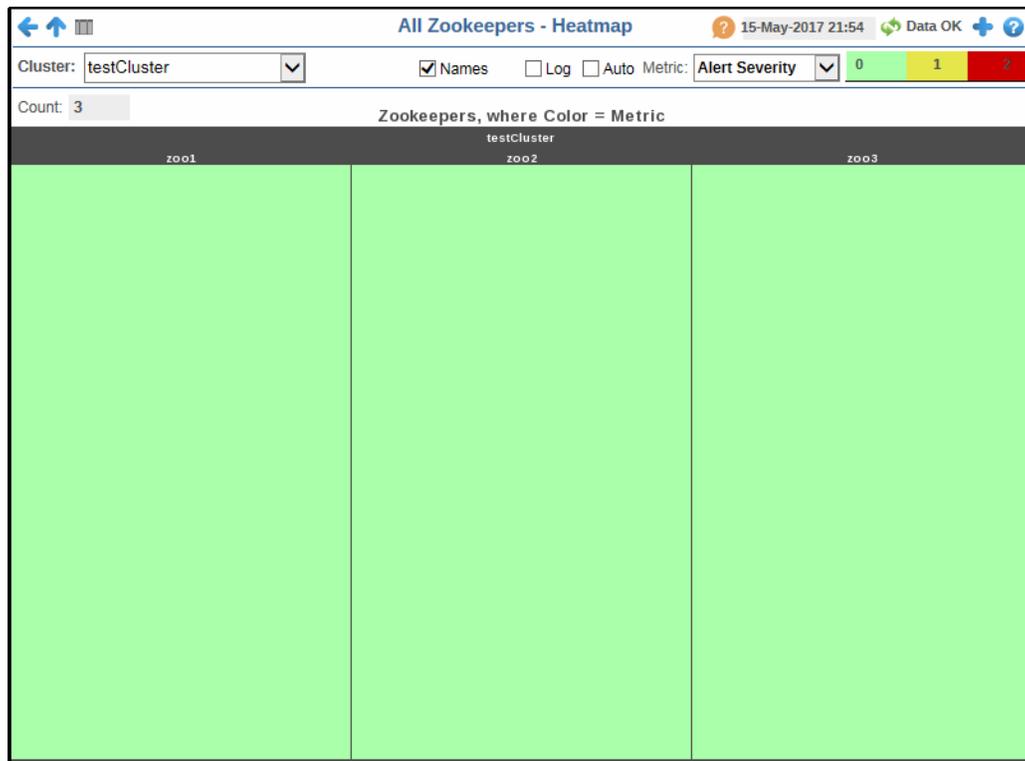
These displays provide detailed data for all zookeepers or for a particular zookeeper. The available displays in this View are:

- ["All Zookeepers Heatmap"](#): Heatmap view of all zookeepers and their associated metrics in a particular cluster.
- ["All Zookeepers Table"](#): Tabular view of all zookeepers and their associated metrics in a particular cluster.
- ["Zookeepers Summary"](#): Contains current and historical metrics, as well as trend data, for a single zookeeper.
- 

### All Zookeepers Heatmap

This heatmap display provides an easy-to-view interface that allows you to quickly identify the current status of each of your zookeepers for each available metric. You can view the zookeepers in the heatmap based on the following metrics: the current alert severity, the current alert count, the number of clients connections, the number of queued requests, the number of incoming packets per second, and the number of outgoing packets per second. By default, this display shows the heatmap based on the **Alert Severity** metric.

You can use the **Names** check-box  to include or exclude labels in the heatmap, and you can mouse over a rectangle to see additional metrics for a zookeeper. Clicking one of the rectangles in the heatmap opens the ["Zookeepers Summary"](#) display, which allows you to see additional details for the selected zookeeper.



**Title Bar** (possible features are):

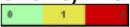
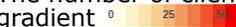
- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Filter By:**

- Cluster** Select the cluster for which you want to view data.
- Names** Select this check box to display the names of the zookeepers at the top of each rectangle in the heatmap.
- Log** Select this check box to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Auto** Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value.  
**Note:** Some metrics auto-scale automatically, even when **Auto** is not selected.
- Metric** Choose a metric to view in the display.

<b>Alert Severity</b>	<p>The current alert severity. Values range from <b>0</b> - <b>2</b>, as indicated in the color gradient  bar, where <b>2</b> is the highest Alert Severity:</p> <ul style="list-style-type: none"> <li> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Alert Count</b>	<p>The total number of critical and warning unacknowledged alerts in the adapters. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.</p>
<b># Alive Connections</b>	<p>The number of clients connected to the zookeeper. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of <b>KafkaZookeeperNumAliveConns</b>. The middle value in the gradient bar indicates the middle value of the range.</p> <p>When <b>Auto</b> is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>
<b>Outstanding Reqs</b>	<p>The number of queued requests. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of <b>KafkaZookeeperOutstandingReqs</b>. The middle value in the gradient bar indicates the middle value of the range.</p> <p>When <b>Auto</b> is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>
<b>Packets In Per Sec</b>	<p>The rate of incoming packets (per second). The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of <b>KafkaZookeeperRatePktsRcvd</b>. The middle value in the gradient bar indicates the middle value of the range.</p> <p>When <b>Auto</b> is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>
<b>Packets Out Per Sec</b>	<p>The rate of outgoing packets (per second). The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of <b>KafkaZookeeperRatePktsSent</b>. The middle value in the gradient bar indicates the middle value of the range.</p> <p>When <b>Auto</b> is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>

## All Zookeepers Table

The table in this display provides a view of all of the zookeepers for a specific cluster and their associated metric data including connection, cluster name, alert level, alert count, and the current value of each gathered metric. You can click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by clicking a row to view details for the selected adapter in the “Zookeepers Summary” display.

Cluster Name	Zookeeper Name	Alert Level	Alert Count	Role	Max Request Latency (ms)	Avg Request Latency (ms)
testCluster	zoo1	●	0	Follower	13	
testCluster	zoo2	●	0	Follower	18	
testCluster	zoo3	●	0	Leader	70	

### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the selected cluster. Refer to KAFKA documentation for more information regarding these fields.

### Filter By:

**Cluster** Select the cluster for which you want to see data.

**Count** The number of zookeepers that were found in the selected cluster.

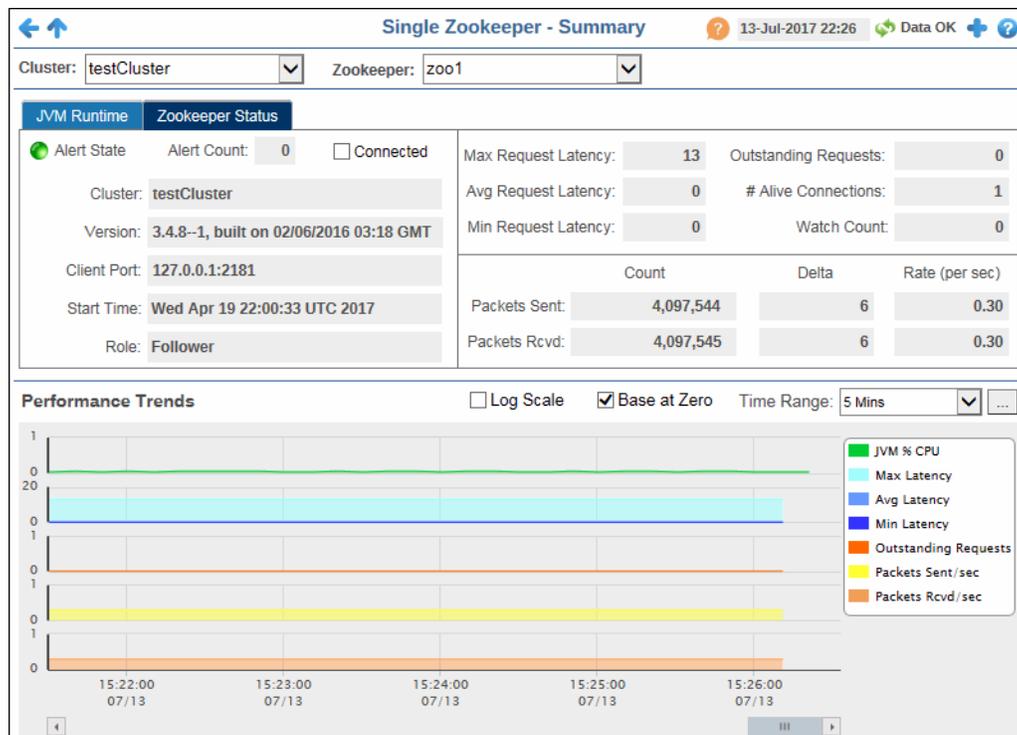
**All Zookeepers Table:**

<b>Cluster Name</b>	The name of the cluster.
<b>Zookeeper Name</b>	The name of the zookeeper.
<b>Alert Level</b>	The current alert severity.  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
<b>Alert Count</b>	The total number of alerts for the host.
<b>Role</b>	The role of the zookeeper ( <b>Leader</b> or <b>Follower</b> ).*
<b>Max Request Latency (ms)</b>	The longest amount of time taken to respond to a client request (in milliseconds) on the zookeeper since the last polling update.*
<b>Avg Request Latency</b>	The average amount of time taken to respond to a client request (in milliseconds) on the zookeeper since the last polling update.*
<b>Min Request Latency</b>	The least amount of time taken to respond to a client request (in milliseconds) on the zookeeper since the last polling update.*
<b>Num Alive Connections</b>	The number of clients connected to the zookeeper.*
<b>Outstanding Requests</b>	The number of queued requests.*
<b>Node Count</b>	The total number of nodes.*
<b>Watch Count</b>	The number of watchers set up over the zookeeper nodes.*
<b>Packets Recvd</b>	The number of packets received.*
<b>Packets Sent</b>	The number of packets sent.*
<b>Delta Packets Recvd</b>	The increase in the amount of packets received by the zookeeper (from the previous polling period to the current polling period).*
<b>Delta Packets Sent</b>	The increase in the amount of packets sent from the zookeeper (from the previous polling period to the current polling period).*
<b>Rate Packets Recvd</b>	The rate at which packets are being received by the zookeeper.*
<b>Rate Packets Sent</b>	The rate at which packets are being sent by the zookeeper.*
<b>Max Client Cnxns Per Host</b>	The maximum number of connections allowed from each host.*
<b>Max Session Timeout</b>	The maximum allowed session timeout allowed for registered consumers.*
<b>Min Session Timeout</b>	The minimum allowed session timeout allowed for registered consumers.*
<b>Version</b>	The current version of Kafka being used.*
<b>Client Port</b>	Lists the client's port.*

<b>JMX Connection String</b>	Lists the connection string.*
<b>Connected</b>	Denotes whether or not the zookeeper is connected.
<b>Expired</b>	When checked, performance data in the row has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the RTView Configuration Application > ( <b>KAFKAMON-LOCAL/Project Name</b> ) > <b>Solution Package Configuration</b> > <b>Apache Kafka</b> > <b>DATA STORAGE</b> > <b>Duration</b> > <b>Expire Time</b> property. The RTView Configuration Application > ( <b>KAFKAMON-LOCAL/Project Name</b> ) > <b>Solution Package Configuration</b> > <b>Apache Kafka</b> > <b>DATA Storage</b> > <b>Duration</b> > <b>Delete Time</b> property allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.  For example, if <b>Expire Time</b> was set to 120 and <b>Delete Time</b> was set to 3600, then the <b>Expired</b> check box would be checked after 120 seconds and the row would be removed from the table after 3600 seconds.
<b>Start Time</b>	The date and time the zookeeper was started.*
<b>Timestamp</b>	The date and time the row data was last updated.

## Zookeepers Summary

This display provides a view of the current and historical metrics for a single zookeeper. You can view JVM runtime statistics and trend data as well as zookeeper status and trend data for the selected zookeeper.



**Title Bar** (possible features are):

-   Open the previous and upper display.
-  Open an instance of this display in a new window.
-  Open the online help page for this display.

open commonly accessed displays.

 **Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

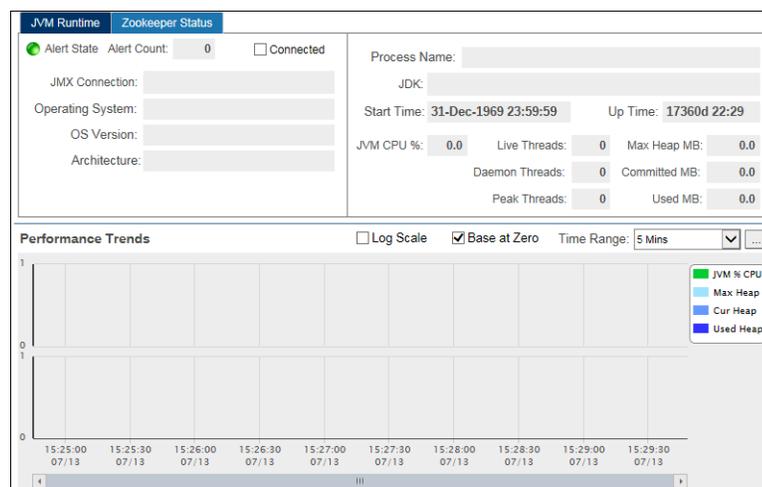
**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the selected cluster. Refer to KAFKA documentation for more information regarding these fields.

### Filter By:

**Cluster** Select the cluster for which you want to show data in the display.

**Zookeeper** Select the zookeeper for which you want to show data in the display.

### JVM Runtime Tab:



**Alert State** The current alert severity.

-  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
-  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
-  Green indicates that no metrics have exceeded their alert thresholds.

**Alert Count** The total number of current alerts.

**Connected** Denotes whether or not the jmx connection is connected.

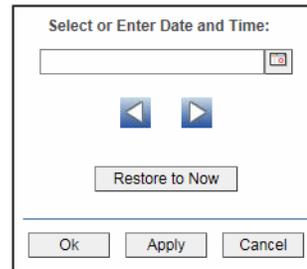
**JMX Connection** The name of the JMX connection.\*

**Operating System** The operating system installed on the zookeeper.\*

<b>OS Version</b>	The version number of the operating systems.*
<b>Architecture</b>	The type of processor being used.*
<b>Process Name</b>	The name of the process.*
<b>JDK</b>	The JDK version number.*
<b>Start Time</b>	The date and time when the zookeeper was started.*
<b>Up Time</b>	The amount of time the zookeeper has been up and running.*
<b>JVM CPU %</b>	The percentage of CPU being used by the JVM.*
<b>Live Threads</b>	The number of live threads.*
<b>Max Heap MB</b>	The maximum amount of available heap, in megabytes.*
<b>Daemon Threads</b>	The number of daemon threads running.*
<b>Committed MB</b>	The total number of megabytes committed.*
<b>Peak Threads</b>	The highest number of threads running at one time during the current polling period.*
<b>Used MB</b>	The number of used megabytes.*
<b>Performance Trends</b>	Traces the following: <ul style="list-style-type: none"> <li><b>JVM % CPU</b> -- traces the percentage of CPU being used by the JVM.</li> <li><b>Max Heap</b> -- traces the maximum amount of available heap.</li> <li><b>Cur Heap</b>-- traces the current amount of heap being used.</li> <li><b>Used Heap</b>-- traces the highest amount of heap used.</li> </ul>
<b>Log Scale</b>	Select to enable a logarithmic scale. Use <b>Log Scale</b> to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. <b>Log Scale</b> makes data on both scales visible by applying logarithmic values rather than actual values to the data.
<b>Base at Zero</b>	Select to use zero ( <b>0</b> ) as the Y axis minimum for all graph traces.

**Time Range**

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .

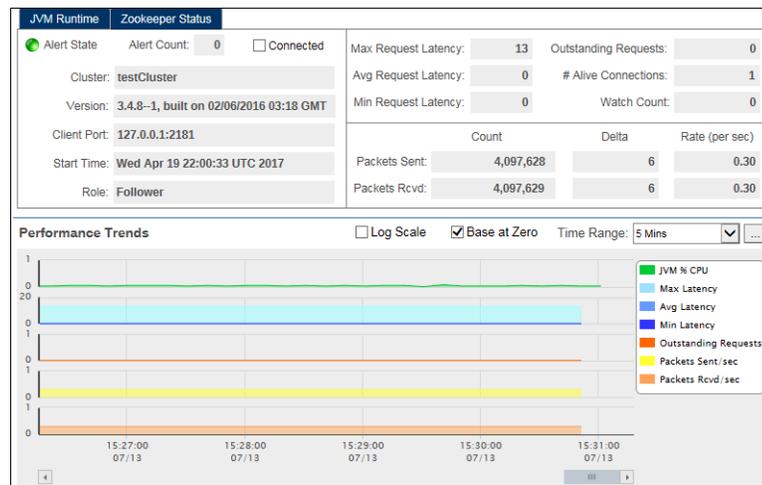


The dialog box titled "Select or Enter Date and Time:" contains a text input field with a calendar icon on the right. Below the input field are two navigation arrows (left and right). At the bottom of the dialog are three buttons: "Restore to Now", "Ok", "Apply", and "Cancel".

By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

**Zookeeper Status Tab:****Alert State**

The current alert severity.

-  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
-  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
-  Green indicates that no metrics have exceeded their alert thresholds.

**Alert Count**

The total number of current alerts.

**Connected**

When checked, denotes that the zookeeper is connected.

**Cluster**

The name of the cluster in which the zookeeper is contained.

**Version**

The current version of Apache Kafka installed.\*

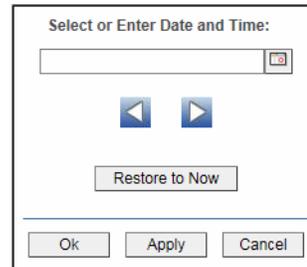
**Client Port**

The client's IP address and port.\*

<b>Start Time</b>	The date and time when the zookeeper was started.*
<b>Role</b>	The zookeeper's role ( <b>Leader/Follower</b> ).*
<b>Max Request Latency</b>	The longest amount of time taken to respond to a client request (in milliseconds) on the zookeeper since the last polling update.*
<b>Avg Request Latency</b>	The average amount of time taken to respond to a client request (in milliseconds) on the zookeeper since the last polling update.*
<b>Min Request Latency</b>	The least amount of time taken to respond to a client request (in milliseconds) on the zookeeper since the last polling update.*
<b>Outstanding Requests</b>	The number of queued requests.*
<b># Alive Connections</b>	The number of clients connected to the zookeeper.*
<b>Watch Count</b>	The number of watchers set up over the zookeeper nodes.*
<b>Packets Sent</b>	<p><b>Count</b> -- The number of packets sent.*</p> <p><b>Delta</b> -- The increase in the amount of packets sent from the zookeeper (from the previous polling period to the current polling period).*</p> <p><b>Rate (per sec)</b> -- The rate at which packets are being sent (per second) by the zookeeper.*</p>
<b>Packets Rcvd</b>	<p><b>Count</b> -- The number of packets received.*</p> <p><b>Delta</b> -- The increase in the amount of packets received by the zookeeper (from the previous polling period to the current polling period).*</p> <p><b>Rate (per sec)</b> -- The rate at which packets are being received (per second) by the zookeeper.*</p>
<b>Performance Trends</b>	<p>Traces the following:</p> <p><b>JVM % CPU</b> -- traces the percentage of CPU used by the JVM.</p> <p><b>Max Latency</b> -- traces the longest amount of time taken to respond to a client request.</p> <p><b>Avg Latency</b> -- traces the average amount of time taken to respond to a client request.</p> <p><b>Min Latency</b> -- traces the least amount of time taken to respond to a client request.</p> <p><b>Outstanding Requests</b> -- traces the number of queued requests.</p> <p><b>Packets Sent/sec</b> -- traces the rate at which packets are being sent (per second) by the zookeeper.</p> <p><b>Packets Rcvd/sec</b> -- traces the rate at which packets are being received (per second) by the zookeeper.</p> <p><b>Log Scale</b> Select to enable a logarithmic scale. Use <b>Log Scale</b> to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. <b>Log Scale</b> makes data on both scales visible by applying logarithmic values rather than actual values to the data.</p>

**Base at Zero** Select to use zero (**0**) as the Y axis minimum for all graph traces.

**Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Kafka Producers View

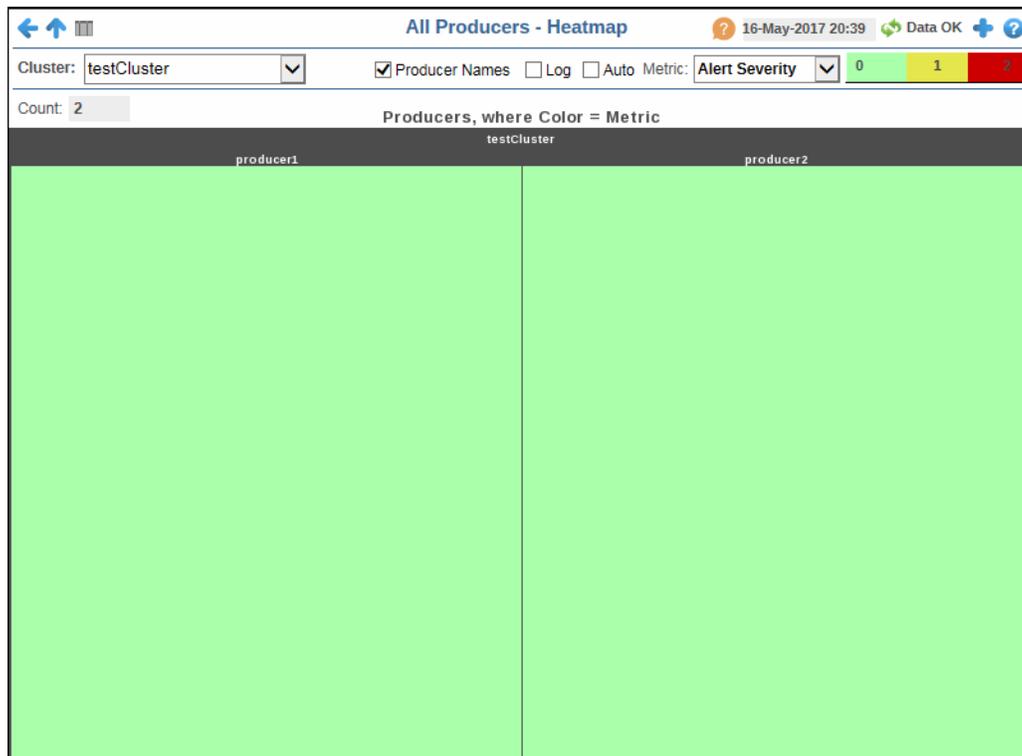
These displays provide detailed data for all producers or for a particular producer. The available displays in this View are:

- **"All Producers Heatmap"**: Heatmap view of all producers and their associated metrics in a particular cluster.
- **"All Producers Table"**: Tabular view of all producers and their associated metrics in a particular cluster.
- **"Producer Summary"**: Contains current and historical metrics, as well as trend data, for a single producer.

### All Producers Heatmap

This heatmap display provides an easy-to-view interface that allows you to quickly identify the current status of each of your producers for each available metric. You can view the producers in the heatmap based on the following metrics: the current alert severity, the current alert count, the incoming/outgoing byte rate, the IO wait time, the request latency, and the request/response rates. By default, this display shows the heatmap based on the **Alert Severity** metric.

You can use the **Producer Names** check-box  to include or exclude labels in the heatmap, and you can mouse over a rectangle to see additional metrics for a producer. Clicking one of the rectangles in the heatmap opens the **"Producer Summary"** display, which allows you to see additional details for the selected producer.



#### Title Bar (possible features are):

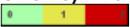
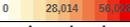
- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

#### Fields and Data:

- Cluster** Select the cluster for which you want to view data.
- Producer Names** Select this check box to display the names of the producers at the top of each rectangle in the heatmap.
- Log** Select this check box to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Auto** Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value.  
**Note:** Some metrics auto-scale automatically, even when **Auto** is not selected.
- Metric** Choose a metric to view in the display.

<b>Alert Severity</b>	<p>The current alert severity. Values range from <b>0</b> - <b>2</b>, as indicated in the color gradient  bar, where <b>2</b> is the highest Alert Severity:</p> <ul style="list-style-type: none"> <li> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Alert Count</b>	<p>The total number of critical and warning unacknowledged alerts in the adapters. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.</p>
<b>Incoming Byte Rate</b>	<p>The rate of incoming bytes (per second). The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of <b>KafkaProducerIncomingByteRate</b>. The middle value in the gradient bar indicates the middle value of the range.</p> <p>When <b>Auto</b> is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>
<b>Outgoing Byte Rate</b>	<p>The rate of outgoing bytes (per second). The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of <b>KafkaProducerOutgoingByteRate</b>. The middle value in the gradient bar indicates the middle value of the range.</p> <p>When <b>Auto</b> is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>
<b>IO Wait Time NSec Avg</b>	<p>The average length of time the IO thread spent waiting for a socket (in nanoseconds). The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of <b>KafkaProducerIoWaitTimeMS</b>. The middle value in the gradient bar indicates the middle value of the range.</p> <p>When <b>Auto</b> is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>
<b>Request Latency</b>	<p>The amount of time between when a producer is called and when the producer receives a response from the broker. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of <b>KafkaProducerRequestLatency</b>. The middle value in the gradient bar indicates the middle value of the range.</p> <p>When <b>Auto</b> is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>

**Request Rate** The average number of requests sent per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **KafkaProducerRequestRate**. The middle value in the gradient bar indicates the middle value of the range.

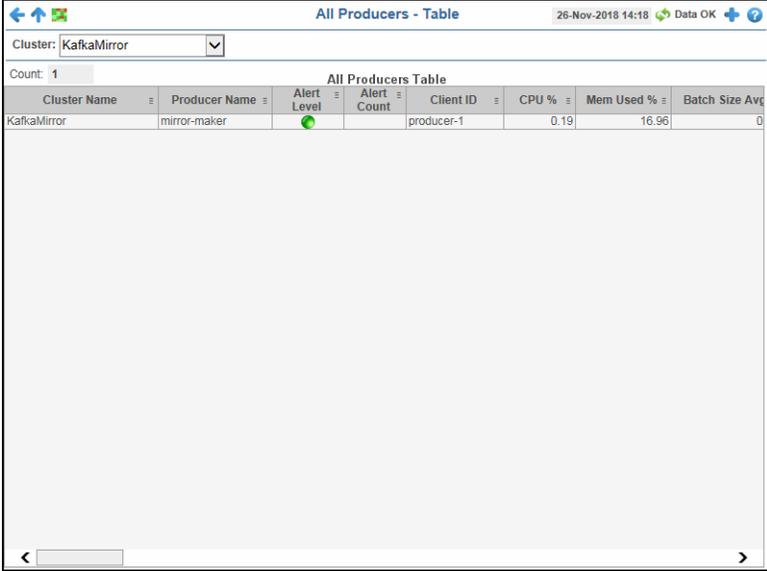
When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

**Response Rate** The average number of responses received (per second). The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **KafkaProducerResponseRate**. The middle value in the gradient bar indicates the middle value of the range.

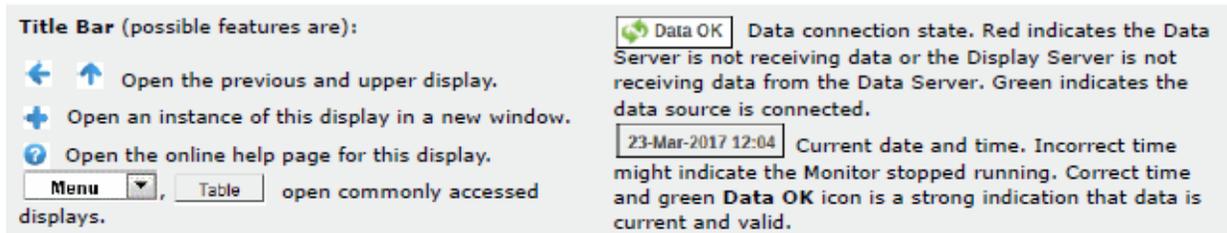
When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

## All Producers Table

The table in this display provides a view of all of your producers and their associated metric data including connection, alert level, alert count, cluster name, client ID, and the current value of each gathered metric. You can click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by clicking a row to view details for the selected adapter in the “[Producer Summary](#)” display.



Cluster Name	Producer Name	Alert Level	Alert Count	Client ID	CPU %	Mem Used %	Batch Size Avg
KafkaMirror	mirror-maker	<span style="color: green;">●</span>		producer-1	0.19	16.96	0



**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the selected cluster. Refer to KAFKA documentation for more information regarding these fields.

### Filter By:

<b>Cluster</b>	Select the cluster for which you want to view data.
<b>Count</b>	The number of producers found on the selected cluster, and that are listed in the <b>All Producers Table</b> .

### All Producers Table:

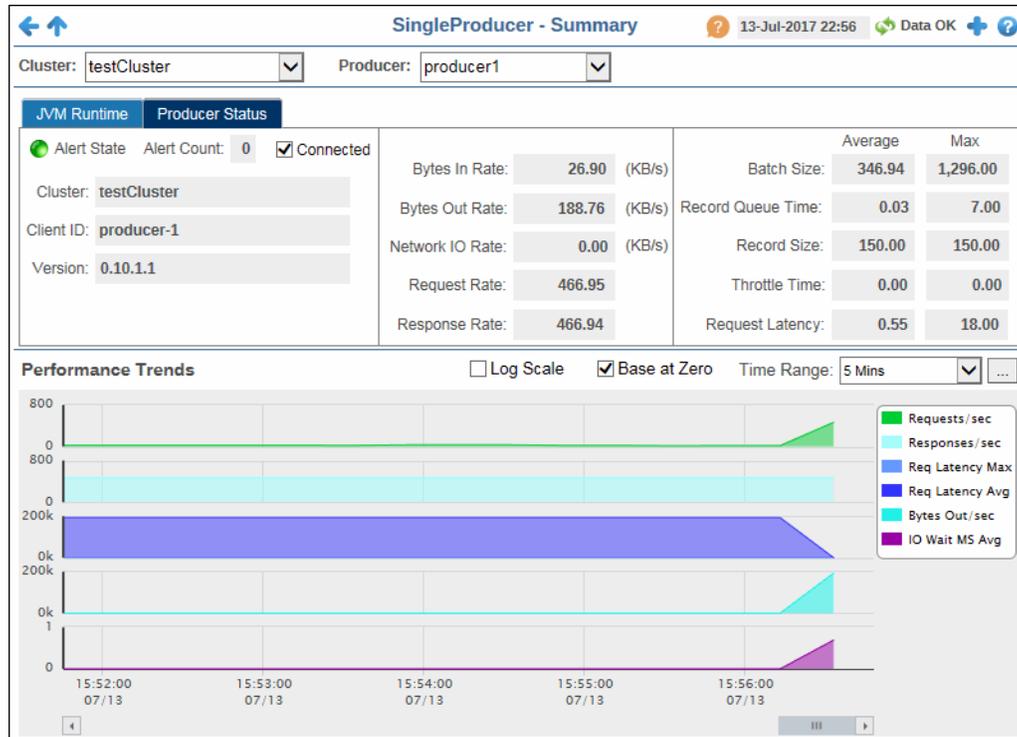
<b>Cluster Name</b>	The name of the cluster.*
<b>Producer Name</b>	The name of the producer.
<b>Alert Level</b>	The current alert severity. <span style="color: red;">●</span> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold. <span style="color: yellow;">●</span> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold. <span style="color: green;">●</span> Green indicates that no metrics have exceeded their alert thresholds.
<b>Alert Count</b>	The total number of alerts for the host.
<b>Client ID</b>	The ID of the producer.*
<b>CPU %</b>	The percentage of CPU being used by the producer.*
<b>Mem Used %</b>	The percentage of JVM memory being used by the producer.*
<b>Batch Size Average</b>	The average batch size sent by the producer.*
<b>Batch Size Max</b>	The maximum number of messages that can be added to a batch before being sent to the event handler.*
<b>Buffer Available Bytes</b>	The number of available bytes in the buffer.*
<b>Buffer Exhausted Rate</b>	The average per-second number of record sends that are dropped due to buffer exhaustion.*
<b>Buffer Total Bytes</b>	The total number of bytes allowed in the buffer.*
<b>Buffer Pool Wait Ratio</b>	The fraction of time an appender waits for space allocation.*

<b>Compression Rate Avg</b>	The average compression rate of record batches.*
<b>Connection Close Rate</b>	The rate of connections being closed.*
<b>Connection Count</b>	The number of active connections.*
<b>Connection Creation Rate</b>	The rate of connections being created.*
<b>Incoming Byte Rate</b>	The average number of incoming bytes per second.*
<b>IO Ratio</b>	The rate of input/output operations.*
<b>IO Time NS Avg</b>	The average length of time the I/O thread spent waiting for a socket (in nanoseconds).*
<b>IO Wait Ratio</b>	The percent of time the producer was performing I/O operations while the CPU was idle.*
<b>IO Wait Time Millisec Avg</b>	The average length of time the I/O thread spent waiting for a socket (in milliseconds).*
<b>Metadata Age</b>	The age (in seconds) of the current producer metadata being used.*
<b>Network IO Rate</b>	The rate of input/output network operations.*
<b>Outgoing Byte Rate</b>	The average number of outgoing bytes per second.*
<b>Produce Throttle Time Avg</b>	The avg time (in milliseconds) a request was throttled by a broker.*
<b>Produce Throttle Time Max</b>	The maximum time (in milliseconds) a request was throttled by a broker.*
<b>Record Error Rate</b>	The average per-second number of record sends that resulted in errors for a topic.*
<b>Record Queue Time Avg</b>	The average time (in milliseconds) record batches spent in the record accumulator.*
<b>Record Queue Time Max</b>	The maximum time (in milliseconds) record batches spent in the record accumulator.*
<b>Record Retry Rate</b>	The average per-second number of retried record sends.
<b>Record Send Rate</b>	The average number of records sent (per second) for a topic.*
<b>Record Size Avg</b>	The average record size.*
<b>Record Size Max</b>	The maximum record size.*
<b>Records per Request Avg</b>	The average number of records per request.*
<b>Request Latency Avg</b>	The average request latency (in milliseconds).*

<b>Request Latency Max</b>	The maximum request latency (in milliseconds).*
<b>Request Rate</b>	The average number of requests sent per second.*
<b>Request Size Avg</b>	The average request size.*
<b>Request Size Max</b>	The maximum request size.*
<b>Requests In Flight</b>	The current number of in-flight requests awaiting a response.*
<b>Response Rate</b>	The average number of responses received per second.*
<b>Select Rate</b>	The number of times the I/O layer checked for new I/O operations to perform per second.*
<b>Waiting Threads</b>	The number of user threads blocked waiting for buffer memory to enqueue their records.*
<b>Jmx Connection String</b>	The JMX connection string.*
<b>Version</b>	The current version of Apache Kafka installed.*
<b>Connected</b>	Denotes whether or not the producer is connected.
<b>Expired</b>	<p>When checked, performance data in the row has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the RTView Configuration Application &gt; (<b>KAFKAMON-LOCAL/Project Name</b>) &gt; <b>Solution Package Configuration</b> &gt; <b>Apache Kafka</b> &gt; <b>DATA STORAGE</b> &gt; <b>Duration</b> &gt; <b>Expire Time</b> property. The RTView Configuration Application &gt; (<b>KAFKAMON-LOCAL/Project Name</b>) &gt; <b>Solution Package Configuration</b> &gt; <b>Apache Kafka</b> &gt; <b>DATA Storage</b> &gt; <b>Duration</b> &gt; <b>Delete Time</b> property allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.</p> <p>For example, if <b>Expire Time</b> was set to 120 and <b>Delete Time</b> was set to 3600, then the <b>Expired</b> check box would be checked after 120 seconds and the row would be removed from the table after 3600 seconds.</p>
<b>Timestamp</b>	The date and time the row data was last updated.

## Producer Summary

This display provides a view of the current and historical metrics for a single producer. You can view JVM runtime statistics and trend data as well as producer status and trend data for the selected producer.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

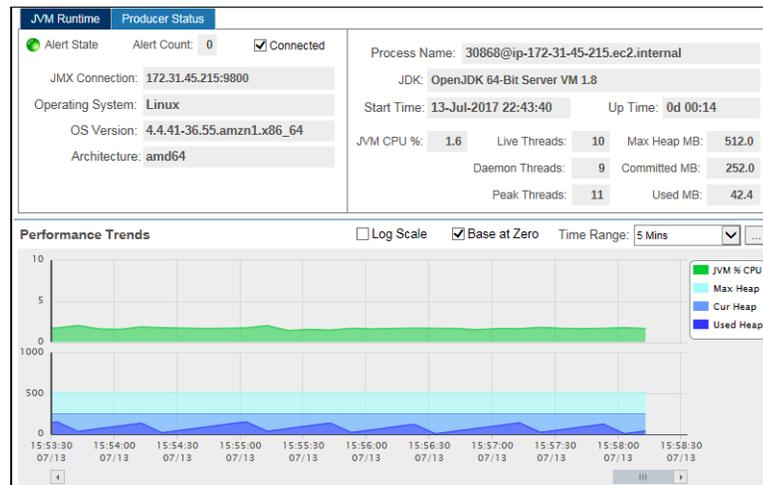
23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

Open the Alert Views - RTView Alerts Table display.

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the selected cluster. Refer to KAFKA documentation for more information regarding these fields.

### Filter By:

**Producer** Select the producer for which you want to show data in the display.

**JVM Runtime Tab:**

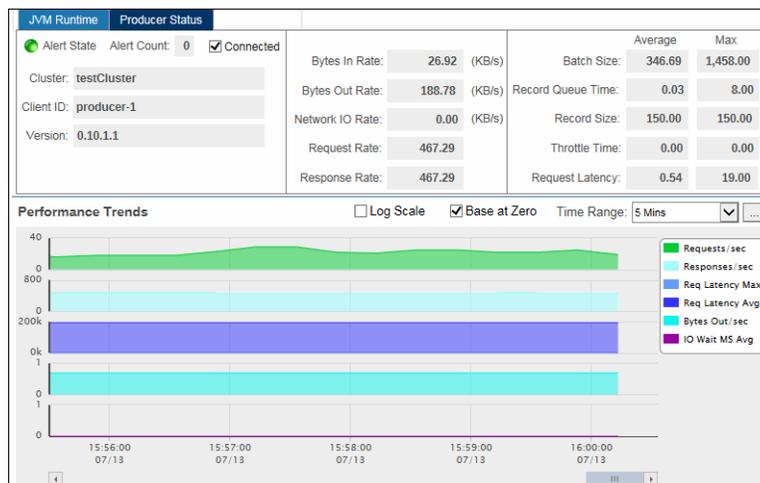
<b>Alert State</b>	The current alert severity. <ul style="list-style-type: none"> <li>● Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li>● Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li>● Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Alert Count</b>	The total number of current alerts.
<b>Connected</b>	Denotes whether or not the jmx connection is connected.
<b>JMX Connection</b>	The name of the JMX connection.*
<b>Operating System</b>	The operating system installed on the producer.*
<b>OS Version</b>	The version number of the operating systems.*
<b>Architecture</b>	The type of processor being used.*
<b>Process Name</b>	The name of the process.*
<b>JDK</b>	The JDK version number.*
<b>Start Time</b>	The date and time when the producer was started.*
<b>Up Time</b>	The amount of time the producer has been up and running.*
<b>JVM CPU %</b>	The percentage of CPU being used by the JVM.*
<b>Live Threads</b>	The number of live threads.*
<b>Max Heap MB</b>	The maximum amount of available heap, in megabytes.*
<b>Daemon Threads</b>	The number of daemon threads running.*
<b>Committed MB</b>	The total number of megabytes committed.*
<b>Peak Threads</b>	The highest number of threads running at one time during the current polling period.*

- Used MB** The number of used megabytes.\*
- Performance Trends** Traces the following:
- JVM % CPU** -- traces the CPU being used by the JVM.
  - Max Heap** -- traces the maximum amount of available heap.
  - Cur Heap**-- traces the current amount of heap being used.
  - Used Heap**-- traces the highest amount of heap used.
- Log Scale** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Base at Zero** Select to use zero (**0**) as the Y axis minimum for all graph traces.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .

By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

**Producer Stats Tab:****Alert State**

The current alert severity.

- Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
- Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
- Green indicates that no metrics have exceeded their alert thresholds.

**Alert Count**

The total number of current alerts.

**Connected**

When checked, denotes that the producer is connected.

**Cluster**

The name of the cluster in which the producer is contained.

**Client ID**

The ID of the client.

**Version**

The current version of Apache Kafka installed.\*

**Bytes In Rate**

The rate of incoming bytes (kilobytes per second).\*

**Bytes Out Rate**

The rate of outgoing bytes (kilobytes per second).\*

**Network IO Rate**

The rate of input/output network operations.\*

**Request Rate**

The average number of requests sent per second.\*

**Response Rate**

The average number of responses received per second.\*

**Batch Size**

**Average** -- The average batch size sent by the producer.\*

**Max** -- The maximum number of messages that can be added to a batch before being sent to the event handler.\*

**Record Queue Time**

**Average** -- The average time (in milliseconds) record batches spent in the record accumulator.\*

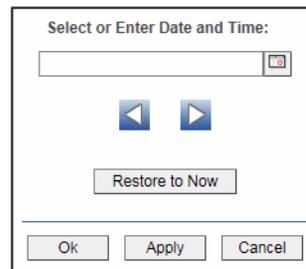
**Max** -- The maximum time (in milliseconds) record batches spent in the record accumulator.\*

**Record Size**

**Average** -- The average record size.\*

**Max** -- The maximum record size.\*

- Throttle Time** **Average** -- The average throttle time (in milliseconds).\*
- Max** -- The maximum time (in milliseconds) a request was throttled by a broker.\*
- Request Latency** **Average** -- The average request latency (in milliseconds).\*
- Max** --The maximum request latency (in milliseconds).\*
- Performance Trends** Traces the following:
- Requests/sec** -- traces the number of requests per second.
  - Responses/sec** -- traces the number of responses per second.
  - Req Latency Max** -- traces the maximum request latency (in milliseconds).
  - Req Latency Avg** -- traces the average request latency (in milliseconds).
  - Bytes Out/sec** -- traces the rate of outgoing bytes (kilobytes per second).
  - IO Wait MS Avg** -- traces the average length of time the I/O thread spent waiting for a socket (in milliseconds).
- Log Scale** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Base at Zero** Select to use zero (0) as the Y axis minimum for all graph traces.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Kafka Consumers View

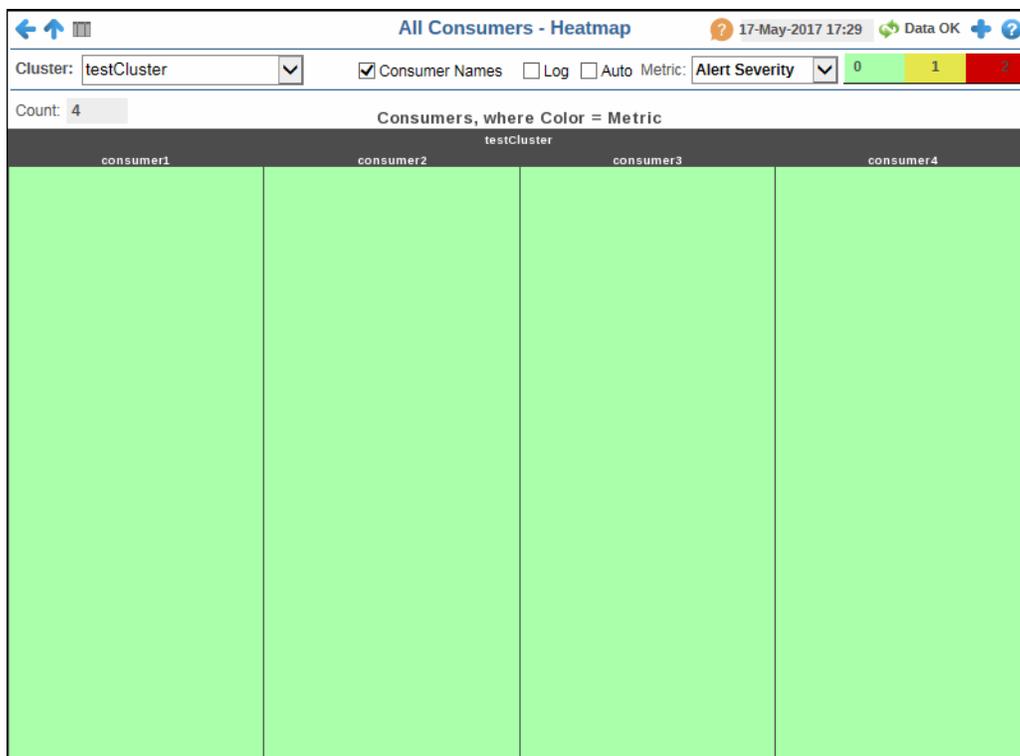
These displays provide detailed data for all consumers or for a particular consumer. The available displays in this View are:

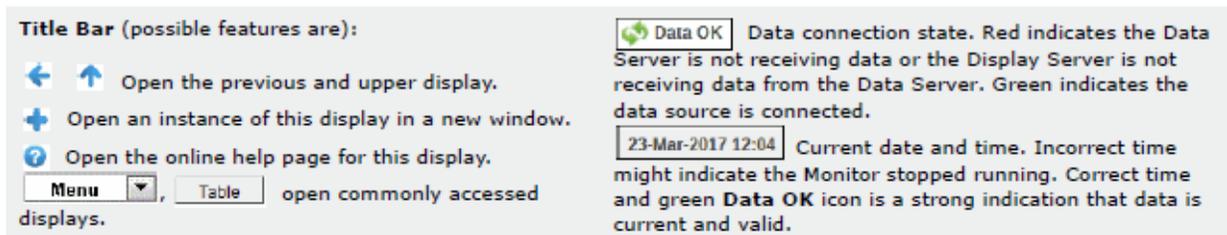
- ["All Consumers Heatmap"](#): Heatmap view of all consumers and their associated metrics in a particular cluster.
- ["All Consumers Table"](#): Tabular view of all consumers and their associated metrics in a particular cluster.
- ["Consumers Summary"](#): Contains current and historical metrics, as well as trend data, for a single consumer.

### All Consumers Heatmap

This heatmap display provides an easy-to-view interface that allows you to quickly identify the current status of each of your consumers for each available metric. You can view the consumers in the heatmap based on the following metrics: the current alert severity, the current alert count, the bytes consumed rate, the fetch latency average, the fetch rate, the maximum consumer lag, and the records consumed rate. By default, this display shows the heatmap based on the **Alert Severity** metric.

You can use the **Consumer Names** check-box  to include or exclude labels in the heatmap, and you can mouse over a rectangle to see additional metrics for a consumer. Clicking one of the rectangles in the heatmap opens the ["Consumers Summary"](#) display, which allows you to see additional details for the selected consumer.





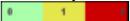
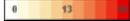
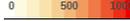
**Title Bar** (possible features are):

-   Open the previous and upper display.
-  Open an instance of this display in a new window.
-  Open the online help page for this display.
-   open commonly accessed displays.

 **Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

## Fields and Data:

- Cluster** Select the cluster for which you want to view data.
- Consumer Names** Select this check box to display the names of the consumers at the top of each rectangle in the heatmap.
- Log** Select this check box to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Auto** Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value.  
**Note:** Some metrics auto-scale automatically, even when **Auto** is not selected.
- Metric** Choose a metric to view in the display.
- Alert Severity** The current alert severity. Values range from **0** - **2**, as indicated in the color gradient  bar, where **2** is the highest Alert Severity:
-  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
  -  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
  -  Green indicates that no metrics have exceeded their alert thresholds.
- Alert Count** The total number of critical and warning unacknowledged alerts in the adapters. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.
- Bytes Consumed Rate** The rate of bytes being consumed (per second). The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **KafkaConsumer**. The middle value in the gradient bar indicates the middle value of the range.
- When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

**Fetch Latency Avg**

The average time taken for fetch request. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **KafkaConsumer**. The middle value in the gradient bar indicates the middle value of the range.

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

**Fetch Rate**

The number of fetch request per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **KafkaConsumer**. The middle value in the gradient bar indicates the middle value of the range.

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

**Consumer Max Lag**

The maximum lag in the number of records for any partition. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **KafkaConsumer**. The middle value in the gradient bar indicates the middle value of the range.

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

**Records Consumed Rate**

The rate of records being consumed (per second). The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **KafkaConsumer**. The middle value in the gradient bar indicates the middle value of the range.

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

## All Consumers Table

The table in this display provides a view of all of your consumers and their associated metric data including connection, alert level, alert count, cluster name, client ID, and the current value of each gathered metric. You can click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by clicking a row to view details for the selected consumer in the ["Consumers Summary"](#) display

Cluster:

Count: 2

Cluster Name	Consumer Name	Alert Level	Alert Count	Client ID	CPU %	Mem Used %	Bytes Consumed
KafkaMirror	mirror-maker	Green		mirror-test-group-0	0.19	15.55	
KafkaMirror	mirror-maker	Green		mirror-test-group-1	0.19	15.55	

**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the selected cluster. Refer to KAFKA documentation for more information regarding these fields.

**Filter By:**

**Cluster** Select the cluster for which you want to view data.

**Count** The number of consumers found on the selected cluster, which are listed in the **All Consumers Table**.

**All Consumers Table:**

**Cluster Name** The name of the cluster.

**Consumer Name** The name of the consumer.

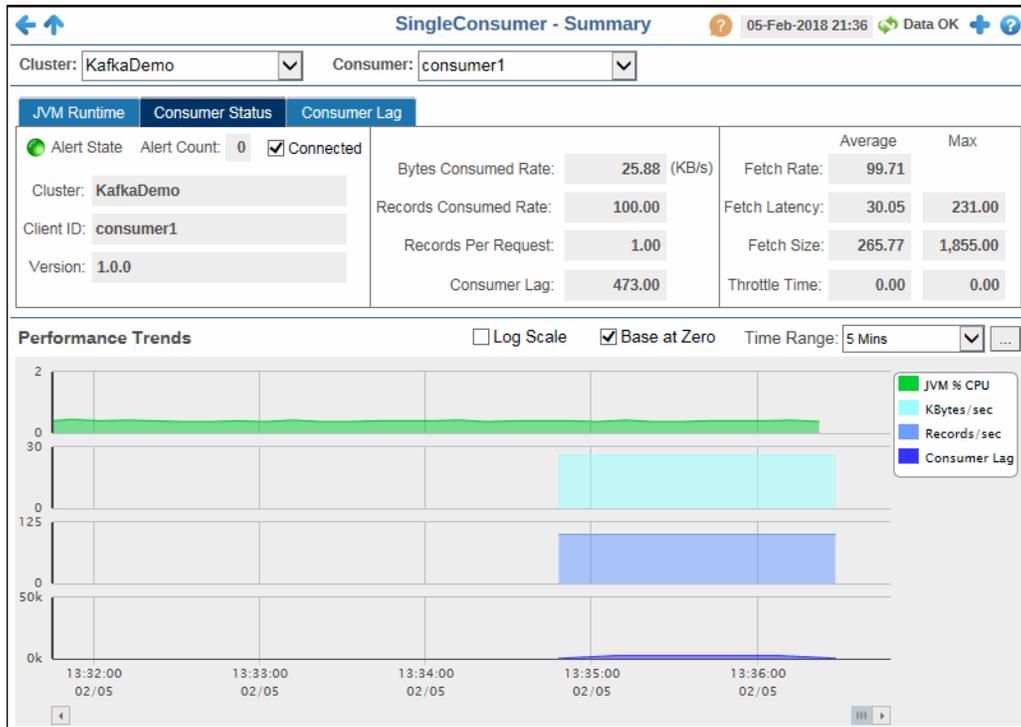
**Alert Level** The current alert severity.

- Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
- Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
- Green indicates that no metrics have exceeded their alert thresholds.

<b>Alert Count</b>	The total number of alerts for the host.
<b>Client ID</b>	The ID of the consumer.*
<b>CPU %</b>	The percentage of CPU being used by the consumer.*
<b>Mem Used %</b>	The percentage of JVM memory being used by the consumer.*
<b>Bytes Consumed Rate</b>	The average number of bytes consumed per second.*
<b>Fetch Latency Avg</b>	The average time taken for a fetch request.*
<b>Fetch Latency Max</b>	The maximum time taken for a fetch request.*
<b>Fetch Rate</b>	The number of fetch requests per second.*
<b>Fetch Size Avg</b>	The average number of bytes fetched per request.*
<b>Fetch Size Max</b>	The maximum number of bytes fetched per request.*
<b>Fetch Throttle Time Avg</b>	The average throttle time in milliseconds.*
<b>Fetch Throttle Time Max</b>	The maximum throttle time in milliseconds.*
<b>Records Consumed Rate</b>	The average number of records consumed per second.*
<b>Records Lag Max</b>	The maximum lag in the number of records for any partition.*
<b>Records per Request Avg</b>	The average number of records in each request.*
<b>JMX Connection String</b>	The JMX connection string.*
<b>Version</b>	The current version of Apache Kafka installed.*
<b>Connected</b>	Denotes whether or not the consumer is connected.*
<b>Expired</b>	<p>When checked, performance data in the row has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the RTView Configuration Application &gt; (<b>KAFKAMON-LOCAL/Project Name</b>) &gt; <b>Solution Package Configuration</b> &gt; <b>Apache Kafka</b> &gt; <b>DATA STORAGE</b> &gt; <b>Duration</b> &gt; <b>Expire Time</b> property. The RTView Configuration Application &gt; (<b>KAFKAMON-LOCAL/Project Name</b>) &gt; <b>Solution Package Configuration</b> &gt; <b>Apache Kafka</b> &gt; <b>DATA Storage</b> &gt; <b>Duration</b> &gt; <b>Delete Time</b> property allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.</p> <p>For example, if <b>Expire Time</b> was set to 120 and <b>Delete Time</b> was set to 3600, then the <b>Expired</b> check box would be checked after 120 seconds and the row would be removed from the table after 3600 seconds.</p>
<b>Timestamp</b>	The date and time the row data was last updated.

## Consumers Summary

This display provides a view of the current and historical metrics for a single consumer. You can view JVM runtime statistics and trend data as well as consumer statistics and trend data for the selected consumer.



**Title Bar** (possible features are):

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu Table open commonly accessed displays.

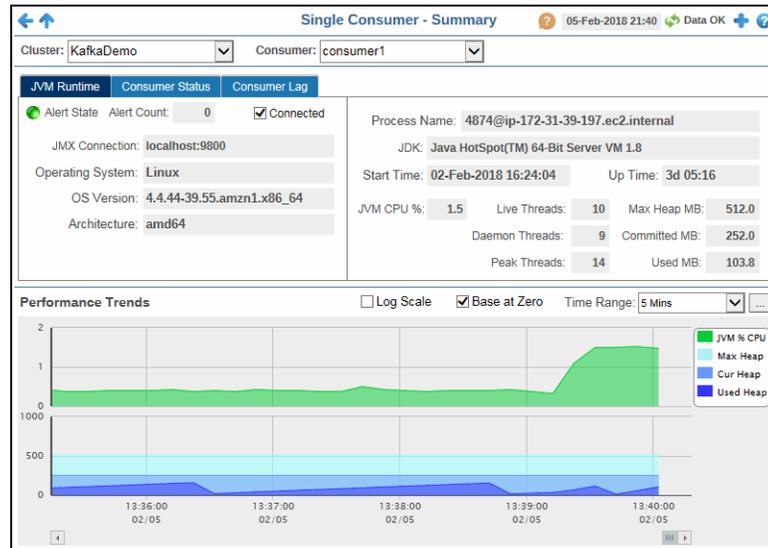
Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the selected cluster. Refer to KAFKA documentation for more information regarding these fields.

**Filter By:**

- Cluster** Select the cluster for which you want to show data in the display.
- Consumer** Select the consumer for which you want to show data in the display.

**JVM Runtime Tab:**

<b>Alert State</b>	The current alert severity. <ul style="list-style-type: none"> <li>● Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li>● Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li>● Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Alert Count</b>	The total number of current alerts.
<b>Connected</b>	Denotes whether or not the JMX connection is connected.
<b>JMX Connection</b>	The name of the JMX connection.*
<b>Operating System</b>	The operating system installed on the producer.*
<b>OS Version</b>	The version number of the operating systems.*
<b>Architecture</b>	The type of processor being used.*
<b>Process Name</b>	The name of the process.*
<b>JDK</b>	The JDK version number.*
<b>Start Time</b>	The date and time when the producer was started.*
<b>Up Time</b>	The amount of time the producer has been up and running.*
<b>JVM CPU %</b>	The percentage of CPU used by the JVM.*
<b>Live Threads</b>	The number of live threads.*
<b>Max Heap MB</b>	The maximum amount of available heap, in megabytes.*
<b>Daemon Threads</b>	The number of daemon threads running.*
<b>Committed MB</b>	The total number of megabytes committed.*

**Peak Threads** The highest number of threads running at one time during the current polling period.\*

**Used MB** The number of used megabytes.\*

**Performance Trends** Traces the following:

**JVM % CPU** -- traces the CPU percentage being used by the JVM.

**Max Heap** -- traces the maximum amount of available heap.

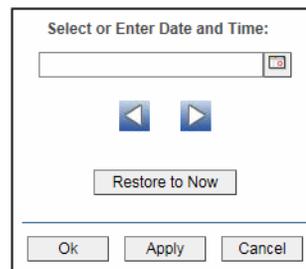
**Cur Heap**-- traces the current amount of heap being used.

**Used Heap**-- traces the highest amount of heap used.

**Log Scale** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.

**Base at Zero** Select to use zero (0) as the Y axis minimum for all graph traces.

**Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .

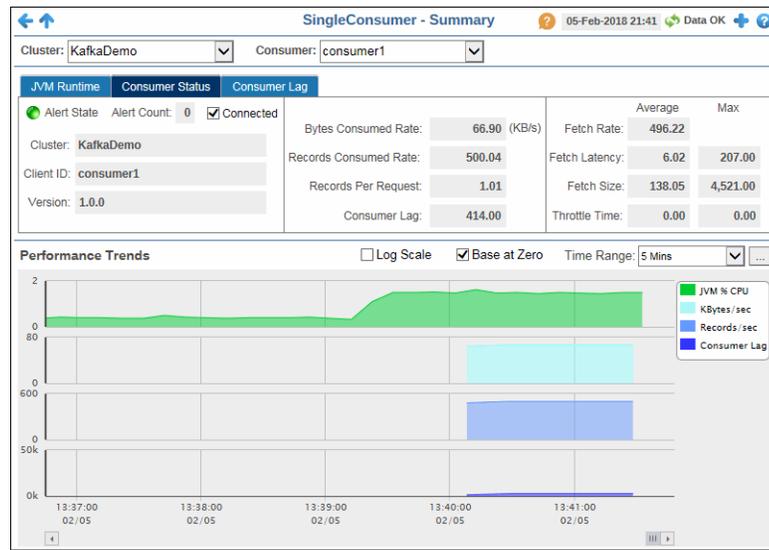


By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Consumer Status Tab:



- Alert State** The current alert severity.
- Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
  - Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
  - Green indicates that no metrics have exceeded their alert thresholds.
- Alert Count** The total number of current alerts.
- Connected** When checked, denotes that the consumer is connected.
- Cluster** The name of the cluster in which the consumer is contained.
- Client ID** The ID of the client.
- Version** The current version of Apache Kafka installed.\*
- Bytes Consumed Rate** The average number of bytes consumed per second.\*
- Records Consumed Rate** The average number of records consumed per second.\*
- Records Per Request** The average number of records in each request.\*
- Consumer Lag** The maximum lag in number of records for any partition.\*
- Fetch Rate** **Average** -- The average number of fetch requests per second.\*  
**Max** -- The highest number of fetch requests per second.\*
- Fetch Latency** **Average** -- The average time taken for a fetch request.\*  
**Max** -- The maximum amount of time taken for a fetch request.\*
- Fetch Size** **Average** -- The average number of bytes fetched per request.\*  
**Max** -- The highest number of bytes fetched per request.\*

**Throttle Time** **Average** -- The average throttle time, in milliseconds.\*  
**Max** -- The maximum throttle time, in milliseconds.\*

**Performance Trends** Traces the following:

**JVM % CPU**-- traces the CPU percentage being used by the JVM.

**KBytes/sec** -- traces the number of kilobytes consumed per second.

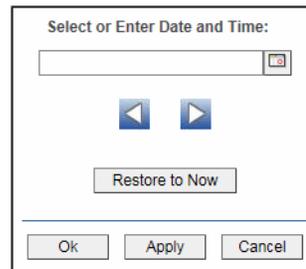
**Records/sec** -- traces the number of records being fetched per second.

**Consumer Lag**-- traces the lag in number of records for any partition.

**Log Scale** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.

**Base at Zero** Select to use zero (0) as the Y axis minimum for all graph traces.

**Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar  .

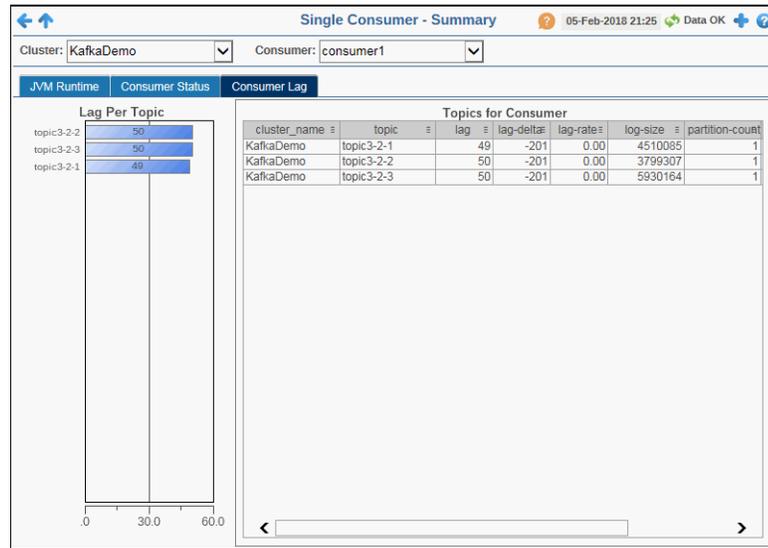


By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Consumer Lag Tab



**Lag Per Topic Bar Graph** Displays the lag per topic in a bar graph format.

### Topics for Consumer Table

<b>cluster_name</b>	The name of the cluster in which the topic resides.
<b>topic</b>	The name of the topic.
<b>lag</b>	The difference between the current consumer position in the partition and the end of the log.*
<b>lag-delta</b>	The difference in the amount of lag from the previous polling period to the current polling period.*
<b>lag-rate</b>	The rate of change in the amount of lag.*
<b>log-size</b>	The current number of messages in the log.*
<b>partition-count</b>	The number of partitions containing the topic.
<b>time_stamp</b>	The date and time the row data was last updated.

---

## Kafka - HTML

The following Kafka Views and their associated displays are available in the Monitor. This section describes the Monitor displays and includes:

- ["Kafka Overview - HTML"](#): Describes the Kafka Overview display.
- ["Kafka Clusters View - HTML"](#): The displays in this View allow you to view metrics for all Kafka clusters and view the performance metrics for all servers on a particular cluster.
- ["Kafka Brokers View - HTML"](#): The displays in this View allow you to view the current and historical metrics for all brokers in heatmap/table formats, view various metrics for a particular broker, and view metrics and trend data for a particular broker.
- ["Kafka Zookeepers View - HTML"](#): The displays in this View allow you to view the current and historical metrics for all zookeepers in a particular cluster in heatmap/tabular format, or view current and historical metrics and trend data for a single zookeeper.
- ["Kafka Topics View - HTML"](#): This displays in this View allow you to view metrics for all topics for a particular broker in heatmap/table format, view current and trend data for a single topic, view the metrics for all topics on a particular cluster, and view metrics for a particular topic on a particular cluster.
- ["Kafka Producers View - HTML"](#): The displays in this View allow you to view the current and historical metrics for all producers in a particular cluster in heatmap/tabular format, or view current and historical metrics and trend data for a single producer.
- ["Kafka Consumers View - HTML"](#): The displays in this View allow you to view the current and historical metrics for all consumers in a particular cluster in heatmap/tabular format, or view current and historical metrics and trend data for a single consumers.

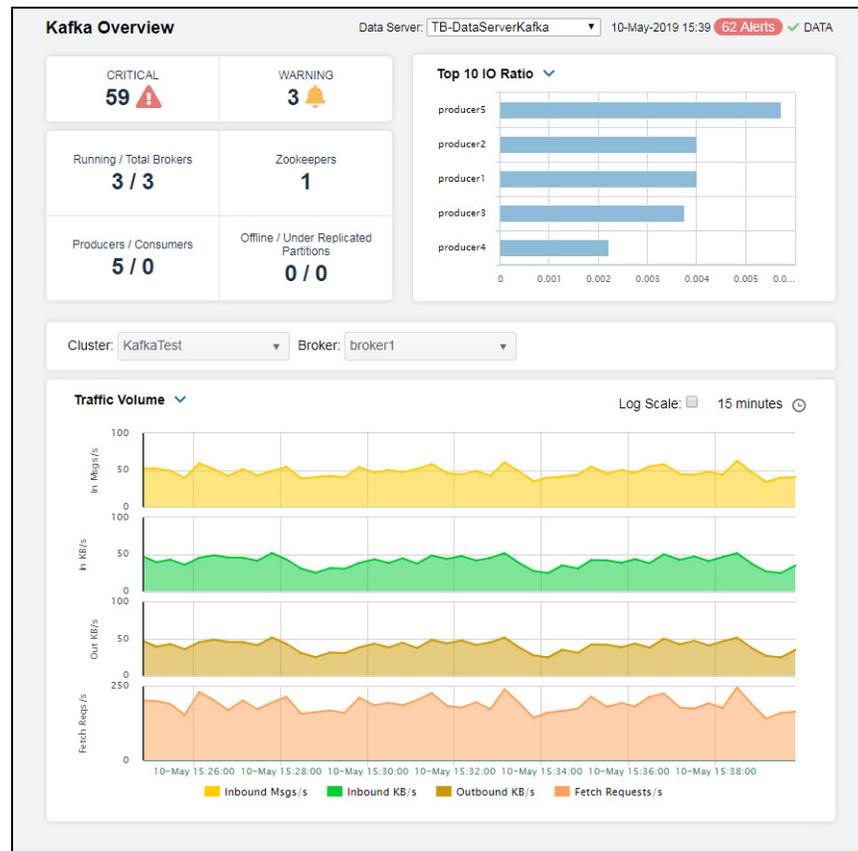
## Kafka Overview - HTML

The **Kafka Overview** is the top-level display for the Kafka Monitor, which provides a good starting point for immediately getting the status of all your clusters, topics, brokers, zookeepers, producers, and consumers on your Data Server. You can select the RTView DataServer for which you want to see data and easily view the current data for that DataServer including:

- The total number of active alerts for the selected DataServer, including the total number of critical and warning alerts.
- The number of running brokers and the total number of brokers.
- The number of zookeepers on your connected DataServer.
- The total number of producers and consumers.
- The number of offline and under-replicated partitions.
- A visual list of the top 10 servers containing the input/output ratio, inbound message rate, consumed kilobytes rate, received packets rate, sent packets rate, and log flush latency (95 percentile) on your connected DataServer.

You can hover over each region in the upper half of the Overview to see more detail. You can also drill down to see even more detail by clicking on each respective region in the Overview. For example, clicking on the alerts in the **CRITICAL** and **WARNING** alerts region opens the **Alerts Table by Components** display.

The bottom half of the display provides a traffic volume and partition availability trend graph for a selected server. You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.



## Kafka Clusters View - HTML

These displays allow you to view metrics for all Kafka clusters and view the performance metrics for all servers on a particular cluster. Clicking **Kafka Clusters** from the left/navigation menu opens the "[Kafka Clusters Table - HTML](#)" display, which shows a tabular view of all clusters and their associated metrics. The option available under **Kafka Clusters** is:

- **Cluster Performance:** Opens the "[Kafka Single Cluster Performance - HTML](#)" display, which allows you to view performance metrics for all servers on a particular cluster.

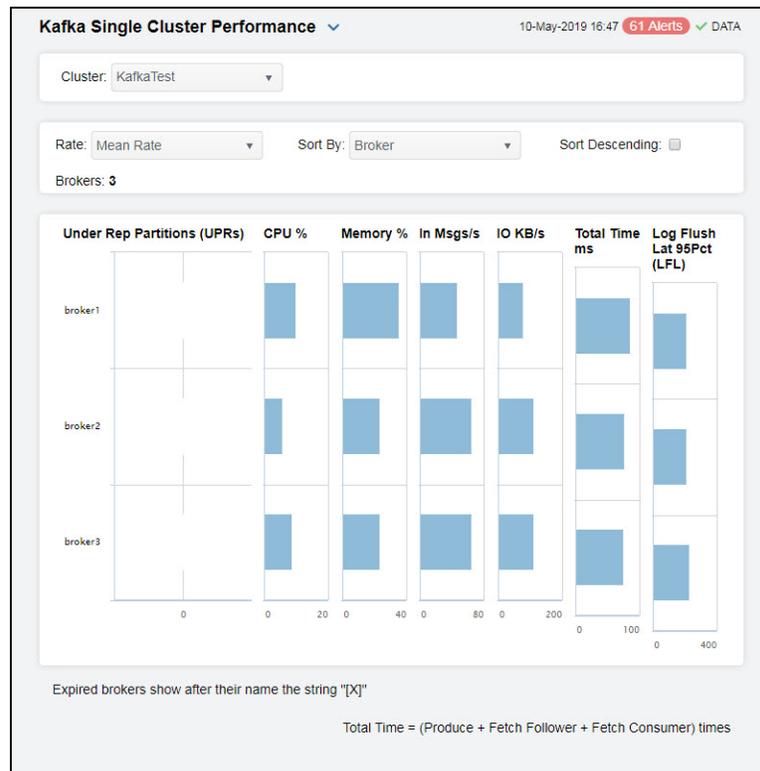
## Kafka Clusters Table - HTML

The **Kafka Clusters Table** contains all metrics available for clusters, including broker, controller, offline partition, and zookeeper data. Each row in the table contains data for a particular cluster. Click a column header to sort column data in ascending or descending order. Double-click on a table row to drill-down to the ["Kafka Single Cluster Performance - HTML"](#) display and view metrics for that particular cluster. Toggle between the commonly accessed displays by clicking the drop down list on the display title.

Cluster	Alert Level	Alert Count	Brokers Monitored	Brokers Running	Active Controllers	Offline Partitions	U
KafkaTest	Warning	61	3	3	1	0	

## Kafka Single Cluster Performance - HTML

Clicking **Cluster Performance** in the left/navigation menu opens the **Kafka Single Cluster Performance** display, which allows you to view of the current metrics for the brokers contained in a selected cluster. The **Rate** drop down list, **Sort By** drop down list, and **Sort Descending** check box all determine the order in which data is displayed in the table. Select an option from the **Rate** drop down list to show data in the table based on the selected rate. Select the metric from the **Sort By** drop down list by which you want to sort the data in the display. When using this option with the **Sort Descending** toggle, the brokers (servers) will be sorted in ascending or descending order using the option you select from this drop down. For example, if you select **In Msgs/s** from this drop down and select the **Sort Descending** toggle, the servers listed in the display will be sorted so that the server with the most **In Msgs/s** will be listed at the top followed by the server with the next most **In Msgs/s**, and so on.



## Kafka Brokers View - HTML

These displays provide detailed data for all brokers in heatmap and tabular form, provide details for all metrics for a particular broker in tabular form, and provide JVM runtime and broker status details for a particular broker. Clicking **Kafka Brokers** from the left/navigation menu opens the "[Kafka Brokers Table - HTML](#)" display, which shows a tabular view of all brokers and their associated metrics. The options available under **Kafka Brokers** are:

- **Brokers Heatmap:** Opens the "[Kafka Brokers Heatmap - HTML](#)" display, which allows you to view performance metrics for all servers on a particular cluster.
- **Single Broker Summary:** Opens the "[Kafka Single Broker Summary - HTML](#)" display, which contains JVM runtime data, broker status, topic, and topic trend details for a particular broker.
- **Single Broker JVM Runtime Summary:** Opens the "[Kafka Single Broker JVM Runtime Summary - HTML](#)" display, which contains JVM runtime data for a single broker.
- **Single Broker Topics Summary:** Opens the "[Kafka Single Broker Topics Summary - HTML](#)" display, which contains topic data for a single broker.
- **Single Broker Topics Lag Summary:** Opens the "[Kafka Single Broker Topic Lag Summary - HTML](#)" display, which displays the lag per topic in a bar graph format and lists the lag per topic for the broker.

## Kafka Brokers Table - HTML

The **Kafka Brokers Table** contains all metrics available for brokers, including partition data, purgatory data, and leader count. Each row in the table contains data for a particular broker. Click a column header to sort column data in ascending or descending order. Double-click on a table row to drill-down to the "[Kafka Single Broker Summary - HTML](#)" display and view metrics for that particular broker. Toggle between the commonly accessed displays by clicking the drop down list on the display title.

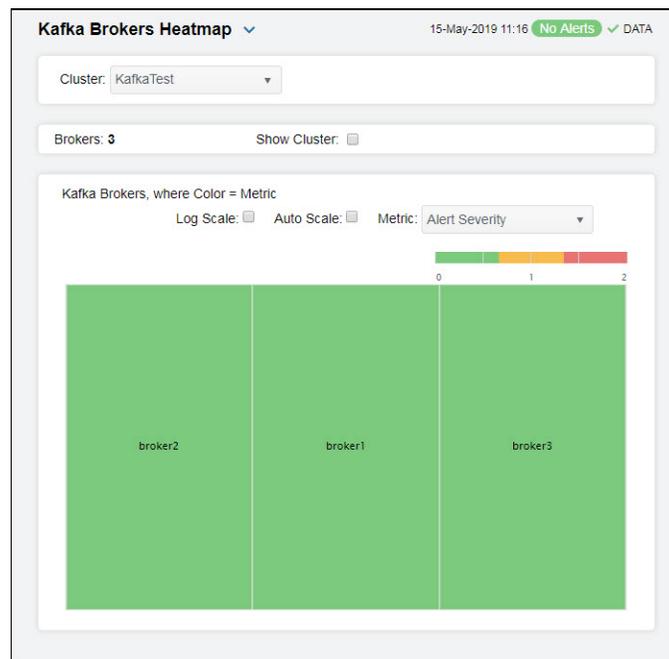
Cluster	Broker	Broker ID	Alert Level	Alert Count	Broker State
KafkaTest	broker1	0	✓		RunningAsBroker
KafkaTest	broker2	1	✓		RunningAsBroker
KafkaTest	broker3	2	✓		RunningAsBroker

## Kafka Brokers Heatmap - HTML

Clicking **Brokers Heatmap** in the left/navigation menu opens the **Kafka Brokers Heatmap**, which allows you to quickly identify the current status of each of your brokers for each available metric. You can view the brokers in the heatmap based on the following metrics: the current alert severity, the current alert count, the under replicated partitions count, the offline partitions count, the rate of incoming messages, the rate of incoming bytes, the rate of outgoing bytes, and the log flush latency value. By default, this display shows the heatmap based on the **Alert Severity** metric.

Each rectangle in the heatmap represents a broker. The rectangle color indicates the most critical alert state associated with the broker. Choose a cluster from the drop-down menu to view all brokers for that cluster. Choose a different metric to display from the **Metric** drop-down menu. Use the **Show Cluster** check-box  to include or exclude labels in the heatmap. Mouse over a rectangle to see additional metrics. By default, this display shows **Alert Severity**.

Drill-down and investigate a broker by clicking a rectangle in the heatmap to view details in the ["Kafka Single Broker Summary - HTML"](#) display.



### Available Metrics

Select the metric driving the heatmap display. The default is **Alert Severity**. Each **Metric** has a color gradient bar that maps values to colors. The heatmap organizes the servers by host, where each rectangle represents a server. Mouse-over any rectangle to display the current values of the metrics for the broker. Click on a rectangle to drill-down to the associated ["Kafka Single Broker Summary - HTML"](#) display for a detailed view of metrics for that particular broker.

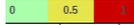
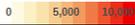
**Alert Severity** The current alert severity. Values range from **0** - **2**, as indicated in the color gradient bar, where **2** is the highest Alert Severity:

- Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
- Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
- Green indicates that no metrics have exceeded their alert thresholds.

**Alert Count** The total number of critical and warning unacknowledged alerts in the brokers. The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.

**Under Replicated Partitions** The number of under-replicated partitions. The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **KafkaBrokerUnderReplicatedPartns**. The middle value in the gradient bar indicates the middle value of the range.

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

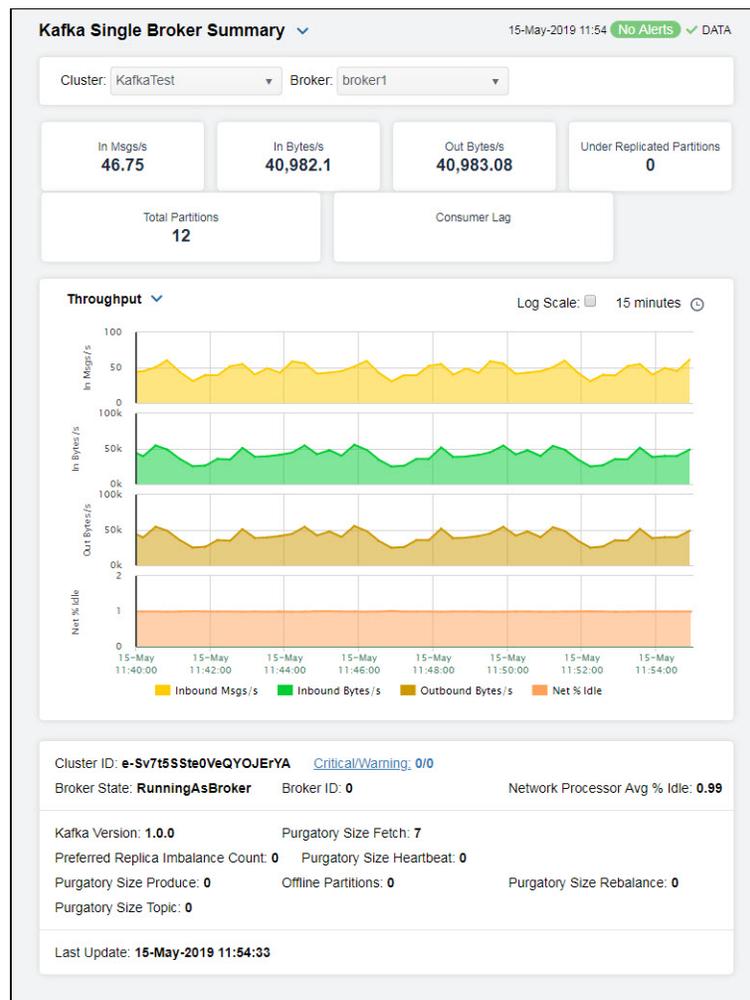
<b>Offline Partitions</b>	<p>The number of offline partitions. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of <b>KafkaBrokerOfflinePartitionCnt</b>. The middle value in the gradient bar indicates the middle value of the range.</p> <p>When <b>Auto</b> is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>
<b>Msgs In Per Sec</b>	<p>The rate of incoming messages (per second). The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of <b>KafkaBrokerMsgsInPerSec</b>. The middle value in the gradient bar indicates the middle value of the range.</p> <p>When <b>Auto</b> is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>
<b>Bytes In Per Sec</b>	<p>The rate of incoming bytes (per second). The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of <b>KafkaBrokerBytesInPerSec</b>. The middle value in the gradient bar indicates the middle value of the range.</p> <p>When <b>Auto</b> is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>
<b>Bytes Out Per Sec</b>	<p>The rate of outgoing bytes (per second). The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of <b>KafkaBrokerBytesOutPerSec</b>. The middle value in the gradient bar indicates the middle value of the range.</p> <p>When <b>Auto</b> is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>
<b>Log Flush Latency 95 Pct</b>	<p>The log flush latency for the top five percent of values. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of <b>KafkaBrokerLogFlushLatency95P</b>. The middle value in the gradient bar indicates the middle value of the range.</p> <p>When <b>Auto</b> is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>

## Kafka Single Broker Summary - HTML

Clicking **Single Broker Summary** in the left/navigation menu opens the **Kafka Single Broker Summary** display, which provides a view of the current and historical metrics for a single broker. Clicking on the information boxes at the top of the display takes you to the "[Kafka Brokers Table - HTML](#)" display, where you can view additional brokers data.

There are two options in the trend graph: **Throughput** and **Partitions**. In the **Throughput** option on the trend graph, you can view trend data for incoming message rate, incoming byte rate, outgoing byte rate, and net percentage idle over a selected time range. In the **Partitions** option on the trend graph, you can view trend data for partitions and active controllers over a selected time range.

Clicking the **Critical/Warning** link at the bottom of the display opens the Alerts Table by Component display.

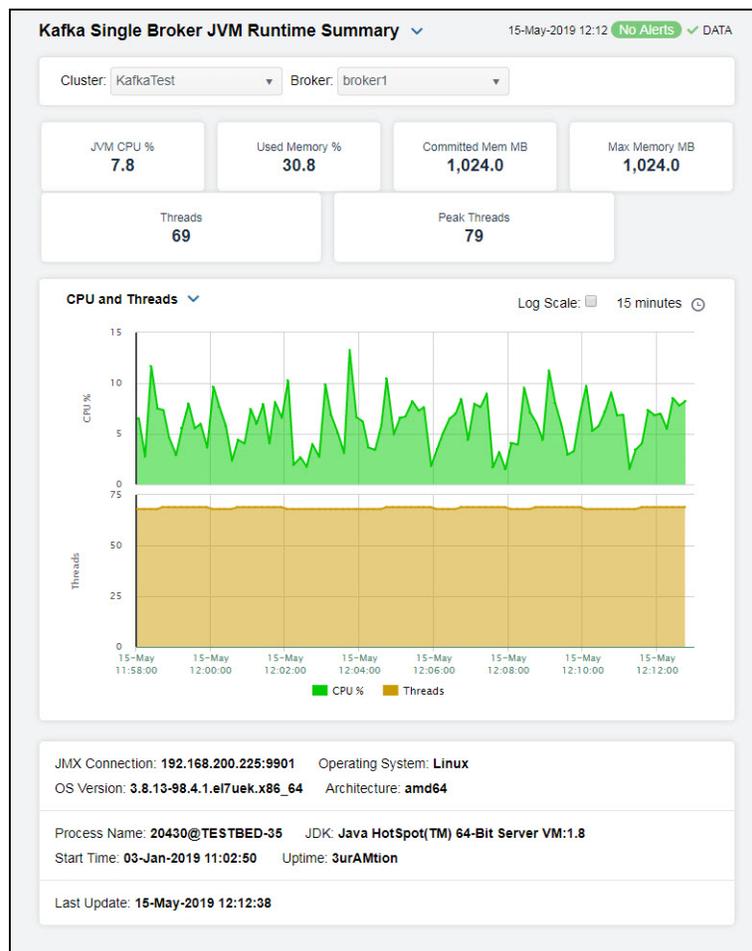


## Kafka Single Broker JVM Runtime Summary - HTML

Clicking **Single Broker JVM Runtime Summary** in the left/navigation menu opens the **Kafka Single Broker JVM Runtime Summary** display, which provides a view of the current and historical JVM Runtime metrics for a single broker. Clicking on the information boxes at the top of the display takes you to the ["Kafka Brokers Table - HTML"](#) display, where you can view additional brokers data.

There are two options in the trend graph: **CPU and Threads** and **Heap Memory**. In the **CPU and Threads** option on the trend graph, you can view trend data for CPU used percentage and number of threads over a selected time range. In the **Heap Memory** option on the trend graph, you can view trend data for the maximum available memory, the used memory, and the committed memory over a selected time range.

Clicking the **Critical/Warning** link at the bottom of the display opens the Alerts Table by Component display.



## Kafka Single Broker Topics Summary - HTML

Clicking **Single Broker Topics Summary** in the left/navigation menu opens the **Kafka Single Broker Topics Summary** display, contains all metrics available for topics for a particular broker. Each row in the table contains data for a particular topic. Click a column header to sort column data in ascending or descending order. Double-click on a table row to drill-down to the "[Kafka Single Topic Summary - HTML](#)" display and view metrics for that particular topic. Toggle between the commonly accessed displays by clicking the drop down list on the display title.

**Kafka Single Broker Topics Summary** ▾
15-May-2019 12:39 ✓ DATA

Cluster: 
Broker:

Rate: 
Topics: 3

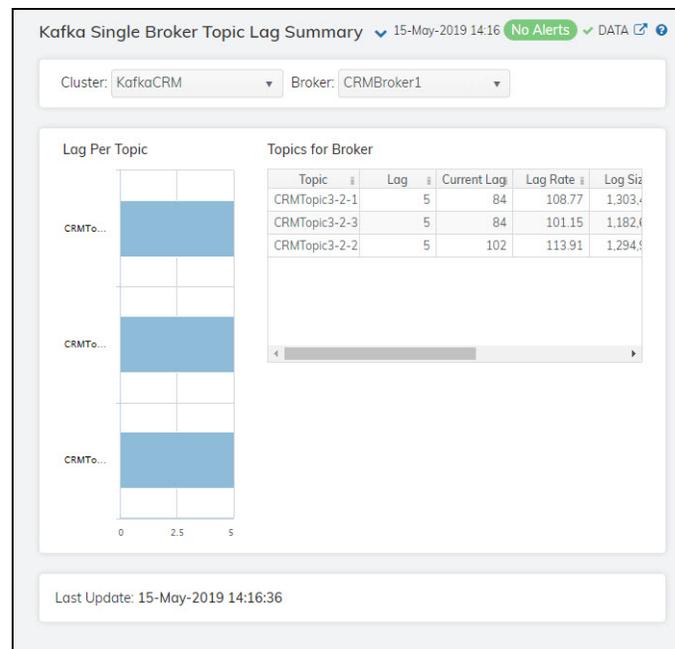
**Metrics by Topic for Selected Broker**

Topic	In Bytes/s	Out Bytes/s	Rejected Bytes/s	Failed Fetched Requests/s	Failed Procs Requests
topic3-2-3	4.338	4.338	0.000	0.000	
topic3-2-2	16413.246	16413.670	0.000	0.000	
topic2-2-1	24565.166	24565.728	0.000	0.000	

## Kafka Single Broker Topic Lag Summary - HTML

Clicking **Single Broker Topics Lag Summary** in the left/navigation menu opens the **Kafka Single Broker Topics Lag Summary** display, which displays the lag per topic in a bar graph format and lists the lag per topic for the broker. Double-click on a bar graph to drill-down to the “[Kafka Single Broker Summary - HTML](#)” display and view metrics for that particular broker.

Each row in the table contains data for a particular topic. Click a column header to sort column data in ascending or descending order. Toggle between the commonly accessed displays by clicking the drop down list on the display title.



## Kafka Zookeepers View - HTML

These displays provide detailed data for all zookeepers or for a particular zookeeper. Clicking **Kafka Zookeepers** from the left/navigation menu opens the “[Kafka Zookeepers Table - HTML](#)” display, which shows a tabular view of all zookeepers and their associated metrics. The options available under **Kafka Zookeepers** are:

- **Zookeepers Heatmap:** Opens the “[Kafka Zookeepers Heatmap - HTML](#)” display, which allows you to view performance metrics for all servers on a particular cluster.
- **Single Zookeeper Summary:** Opens the “[Kafka Single Zookeeper Summary - HTML](#)” display, which contains current and historical metrics, as well as trend data, for a single zookeeper.
- **Single Zookeeper JVM Runtime Summary:** Opens the “[Kafka Single Zookeeper JVM Runtime Summary - HTML](#)” display, which contains JVM Runtime metrics for a single zookeeper.

## Kafka Zookeepers Table - HTML

The table in this display provides a view of all of the zookeepers for a specific cluster and their associated metric data including connection, cluster name, alert level, alert count, and the current value of each gathered metric. Each row in the table contains data for a particular zookeeper. Click a column header to sort column data in ascending or descending order. Double-click on a table row to drill-down to the ["Kafka Single Zookeeper Summary - HTML"](#) display and view metrics for that particular zookeeper. Toggle between the commonly accessed displays by clicking the drop down list on the display title.

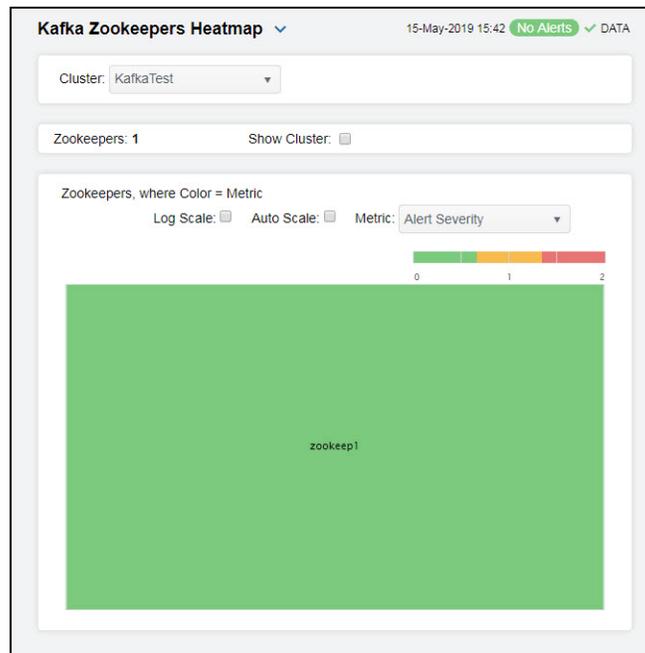
Cluster	Zookeeper	URL	Alert Level	Alert Count	Role
KafkaTest	zookeep1		✓		

## Kafka Zookeepers Heatmap - HTML

Clicking **Zookeepers Heatmap** in the left/navigation menu opens the **Kafka Zookeepers Heatmap**, which allows you to quickly identify the current status of each of your zookeepers for each available metric. You can view the zookeepers in the heatmap based on the following metrics: the current alert severity, the current alert count, the number of clients connections, the number of queued requests, the number of incoming packets per second, and the number of outgoing packets per second. By default, this display shows the heatmap based on the **Alert Severity** metric.

Each rectangle in the heatmap represents a zookeeper. The rectangle color indicates the most critical alert state associated with the zookeeper. Choose a cluster from the drop-down menu to view all zookeepers for that cluster. Choose a different metric to display from the **Metric** drop-down menu. Use the **Show Cluster** check-box  to include or exclude labels in the heatmap. Mouse over a rectangle to see additional metrics. By default, this display shows **Alert Severity**.

Drill-down and investigate a broker by clicking a rectangle in the heatmap to view details in the ["Kafka Single Zookeeper Summary - HTML"](#) display.



## Available Metrics

Select the metric driving the heatmap display. The default is **Alert Severity**. Each **Metric** has a color gradient bar that maps values to colors. The heatmap organizes the servers by host, where each rectangle represents a zookeeper. Mouse-over any rectangle to display the current values of the metrics for the zookeeper. Click on a rectangle to drill-down to the associated ["Kafka Single Zookeeper Summary - HTML"](#) display for a detailed view of metrics for that particular zookeeper.

### Alert Severity

The current alert severity. Values range from **0** - **2**, as indicated in the color gradient  bar, where **2** is the highest Alert Severity:

-  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
-  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
-  Green indicates that no metrics have exceeded their alert thresholds.

### Alert Count

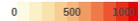
The total number of critical and warning unacknowledged alerts in the adapters. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.

### # Alive Connections

The number of clients connected to the zookeeper. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **KafkaZookeeperNumAliveConns**. The middle value in the gradient bar indicates the middle value of the range.

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

**Outstanding Reqs**

The number of queued requests. The color gradient bar  shows the range of the value/color mapping. The current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **KafkaZookeeperOutstandingReqs**. The middle value in the gradient bar indicates the middle value of the range.

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

**Packets In Per Sec**

The rate of incoming packets (per second). The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **KafkaZookeeperRatePktsRcvd**. The middle value in the gradient bar indicates the middle value of the range.

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

**Packets Out Per Sec**

The rate of outgoing packets (per second). The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **KafkaZookeeperRatePktsSent**. The middle value in the gradient bar indicates the middle value of the range.

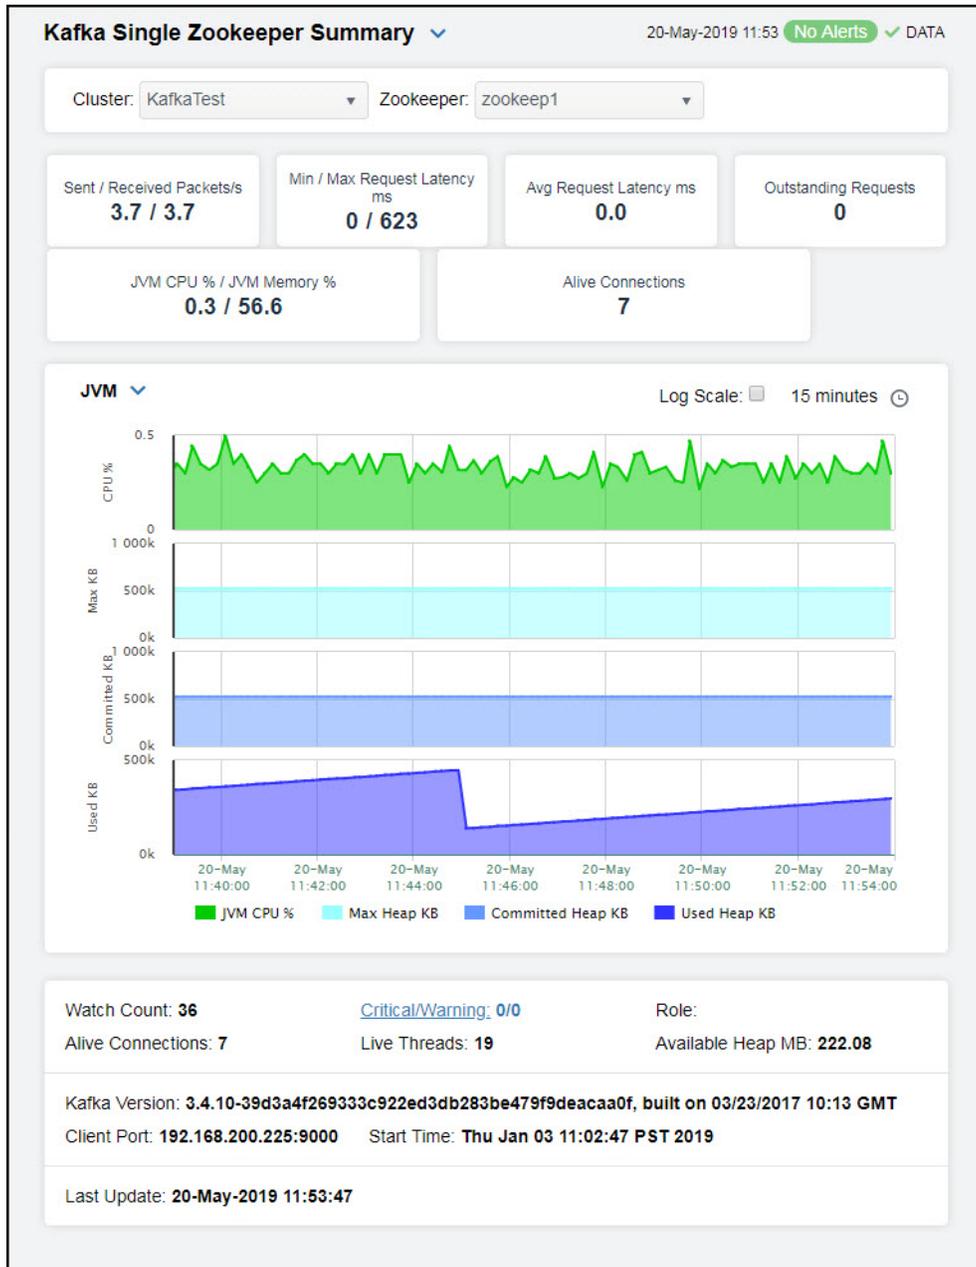
When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

## Kafka Single Zookeeper Summary - HTML

Clicking **Single Zookeeper Summary** in the left/navigation menu opens the **Kafka Single Zookeeper Summary** display, which provides a view of the current and historical metrics for a single zookeeper. Clicking on the information boxes at the top of the display takes you to the ["Kafka Zookeepers Table - HTML"](#) display, where you can view additional zookeepers data.

There are three options in the trend graph: **Latency**, **Performance**, and **JVM**. In the **Latency** option on the trend graph, you can view trend data for maximum latency, average latency, and minimum latency over a selected time range. In the **Performance** option on the trend graph, you can view trend data for the outstanding requests, sent packet rate, received packet rate, and alive connections over a selected time range. In the **JVM** option on the trend graph, you can view trend data for JVM CPU percentage, maximum heap in kilobytes, committed heap in kilobytes, and used heap in kilobytes over a selected time range.

Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display.



## Kafka Single Zookeeper JVM Runtime Summary - HTML

Clicking **Single Zookeeper JVM Runtime Summary** in the left/navigation menu opens the **Kafka Single Zookeeper JVM Runtime Summary** display, which provides a view of the current and historical JVM Runtime metrics for a single zookeeper. Clicking on the information boxes at the top of the display takes you to the "[Kafka Zookeepers Table - HTML](#)" display, where you can view additional zookeepers data.

In the **Performance Trends** trend graph, you can view trend data for CPU percentage, maximum memory in megabytes, committed memory in megabytes, and used memory in megabytes over a selected time range.

Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display.

### Kafka Single Zookeeper Summary ▼ 20-May-2019 11:53 No Alerts DATA

Cluster: KafkaTest Zookeeper: zookeep1

Sent / Received Packets/s <b>3.7 / 3.7</b>	Min / Max Request Latency ms <b>0 / 623</b>	Avg Request Latency ms <b>0.0</b>	Outstanding Requests <b>0</b>
JVM CPU % / JVM Memory % <b>0.3 / 56.6</b>	Alive Connections <b>7</b>		

#### JVM ▼ Log Scale: 15 minutes

Legend: JVM CPU % Max Heap KB Committed Heap KB Used Heap KB

Watch Count: <b>36</b>	<a href="#">Critical/Warning: 0/0</a>	Role:
Alive Connections: <b>7</b>	Live Threads: <b>19</b>	Available Heap MB: <b>222.08</b>

Kafka Version: **3.4.10-39d3a4f269333c922ed3db283be479f9deacaa0f**, built on 03/23/2017 10:13 GMT  
Client Port: **192.168.200.225:9000** Start Time: **Thu Jan 03 11:02:47 PST 2019**

Last Update: **20-May-2019 11:53:47**

## Kafka Topics View - HTML

These displays allow you to view metrics for all Kafka topics on a particular topic in heatmap/table format, view the performance metrics for a single topic on a particular broker, view the metrics for all topics on a particular cluster, and view metrics for a particular topic on a particular cluster. Clicking **Kafka Topics** from the left/navigation menu opens the "[Kafka Topics Table - HTML](#)" display, which shows a tabular view of all clusters and their associated metrics. The options available under **Kafka Topics** are:

- **Kafka Topics Heatmap:** Opens the "[Kafka Topics Heatmap - HTML](#)" display, which allows you to view performance metrics for all servers on a particular cluster.
- **Single Topic Summary:** Opens the "[Kafka Single Topic Summary - HTML](#)" display, which allows you to view current metrics and trend data for a single topic.
- **Single Topic Partition Summary:** Opens the "[Kafka Single Topic Partition Summary - HTML](#)" display, which allows you to view current metrics and trend data for a single topic partition.
- **Topics Activity by Cluster:** Opens the "[Kafka Topics Activity by Cluster - HTML](#)" display, which allows you to view performance metrics for all topics on a particular cluster.
- **Brokers Activity by Topic:** Opens the "[Kafka Brokers Activity by Topic - HTML](#)" display, which allows you to view performance metrics for all brokers for a particular topic.

## Kafka Topics Table - HTML

The table in this display provides a view of all of your topics for a particular cluster and their associated metric data. Each row in the table contains data for a particular topic. Click a column header to sort column data in ascending or descending order. Double-click on a table row to drill-down to the "[Kafka Single Topic Summary - HTML](#)" display and view metrics for that particular topic. Toggle between the commonly accessed displays by clicking the drop down list on the display title.

**Kafka Topics Table** 20-May-2019 14:16 ✓ DATA

Cluster: KafkaTest

Topics: 11      Rate: Mean Rate

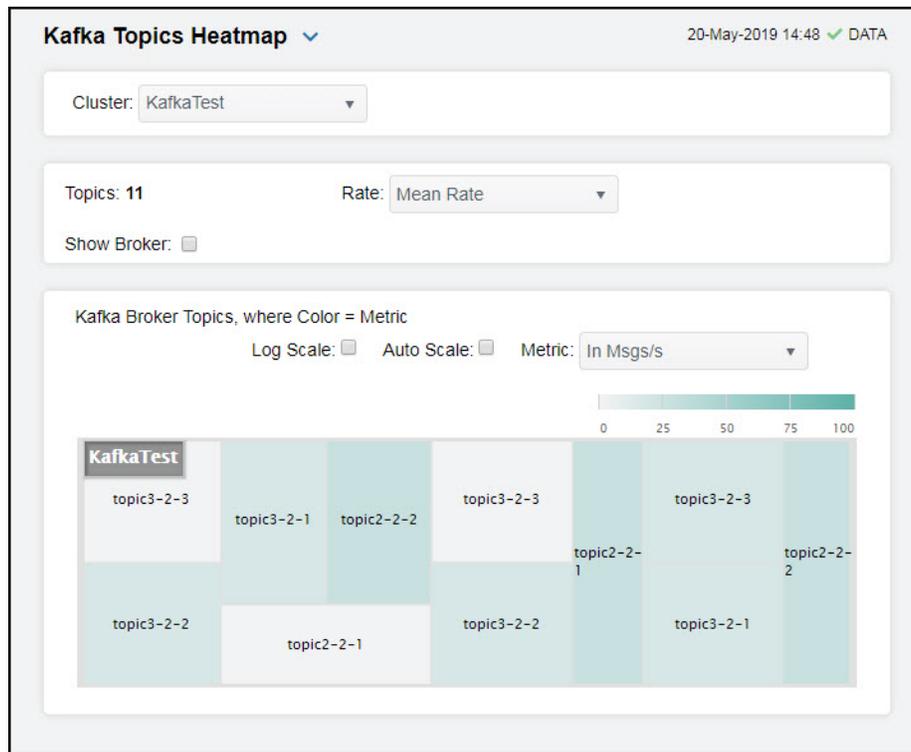
Cluster	Broker	Topic	In Bytes/s	Out Bytes/s	Rejected B
KafkaTest	broker1	topic3-2-3	4.167	4.167	
KafkaTest	broker1	topic3-2-2	16415.206	16415.633	
KafkaTest	broker1	topic2-2-1	24568.572	24569.128	
KafkaTest	broker2	topic3-2-3	4.197	4.197	
KafkaTest	broker2	topic3-2-2	16411.263	16411.615	
KafkaTest	broker2	topic3-2-1	16413.949	16414.330	
KafkaTest	broker2	topic2-2-2	24578.163	24578.852	
KafkaTest	broker2	topic2-2-1	5.278	5.278	
KafkaTest	broker3	topic3-2-3	16413.636	16413.961	
KafkaTest	broker3	topic3-2-1	16413.890	16414.319	
KafkaTest	broker3	topic2-2-2	24578.245	24578.864	

## Kafka Topics Heatmap - HTML

Clicking **Kafka Topics Heatmap** in the left/navigation menu opens the **Kafka Topics Heatmap**, which provides an easy-to-view interface that allows you to quickly identify the current status of each of your topics for each available metric. You can view the topics in the heatmap based on the following metrics: the rate of incoming messages, the rate of incoming bytes, the rate of outgoing bytes, the rate of rejected bytes, the rate of total fetch requests, the rate of failed fetch requests, and the rate of total produce requests. By default, this display shows the heatmap based on the **In Msgs/s** metric.

Each rectangle in the heatmap represents a topic. The rectangle color indicates the most critical alert state associated with the topic. Choose a cluster from the drop-down menu to view all topics for that cluster. Choose a different metric to display from the **Metric** drop-down menu. Use the **Show Broker** check-box  to include or exclude labels in the heatmap. Mouse over a rectangle to see additional metrics.

Drill-down and investigate a topic by clicking a rectangle in the heatmap to view details in the ["Kafka Single Topic Summary - HTML"](#) display.



## Available Metrics

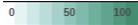
Select the metric driving the heatmap display. The default is **In Msgs/s**. Each **Metric** has a color gradient bar that maps values to colors. Mouse-over any rectangle to display the current values of the metrics for the topic. Click on a rectangle to drill-down to the associated ["Kafka Single Topic Summary - HTML"](#) display for a detailed view of metrics for that particular topic.

- In Msgs/s** The rate of incoming messages (per second). The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum number of incoming messages per second. The middle value in the gradient bar indicates the middle value of the range.
- In Bytes/s** The rate of incoming bytes (per second). The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum number of incoming bytes per second. The middle value in the gradient bar indicates the middle value of the range.
- Out Bytes/s** The rate of outgoing bytes (per second). The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum number of outgoing bytes per second. The middle value in the gradient bar indicates the middle value of the range.
- Rejected Bytes/s** The rate of bytes being rejected (per second). The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum number of bytes rejected per second. The middle value in the gradient bar indicates the middle value of the range.

**Total Fetch Requests/s**

The rate of fetch requests (per second). The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum number of fetch requests per second. The middle value in the gradient bar indicates the middle value of the range.

**Failed Fetch Requests/s**

The rate of failed fetch requests (per second). The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum number of failed fetch requests per second. The middle value in the gradient bar indicates the middle value of the range.

**Total Produce Requests/s**

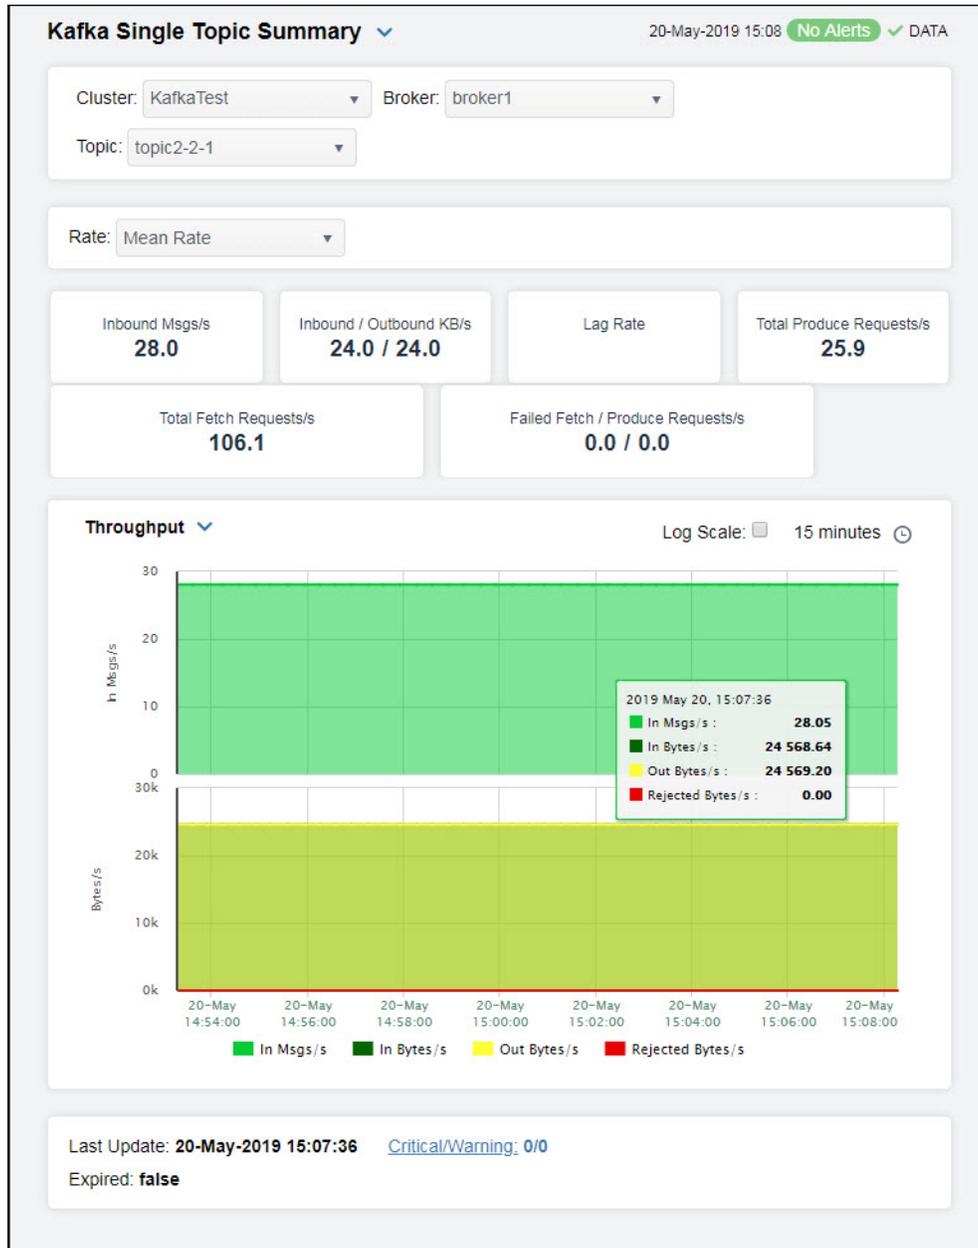
The rate of total producer requests (per second). The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum number of produce requests per second. The middle value in the gradient bar indicates the middle value of the range.

## Kafka Single Topic Summary - HTML

Clicking **Single Topic Summary** in the left/navigation menu opens the **Kafka Single Topic Summary** display, which provides a view of the current metrics and trend data for a single topic. Clicking on the information boxes at the top of the display takes you to the "[Kafka Topics Table - HTML](#)" display, where you can view additional topic data.

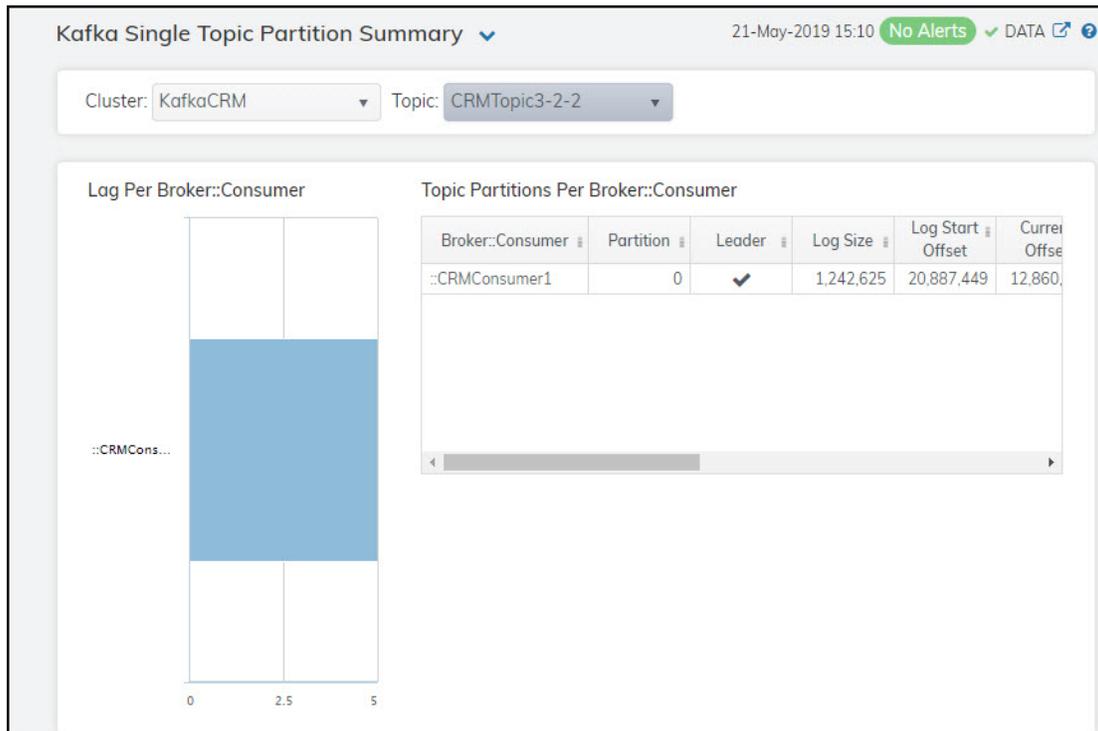
There are three options in the trend graph: **Throughput**, **Produce Requests**, and **Fetch Requests**. In the **Throughput** option on the trend graph, you can view trend data for incoming message rate, incoming byte rate, outgoing byte rate, and rejected byte rate over a selected time range. In the **Produce Requests** option on the trend graph, you can view trend data for the total produce requests rate, failed produce requests rate, and lag rate over a selected time range. In the **Fetch Requests** option on the trend graph, you can view trend data for total fetch requests rate, failed fetch requests rate, and lag rate over a selected time range.

Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display.



## Kafka Single Topic Partition Summary - HTML

Clicking **Single Topic Partition Summary** in the left/navigation menu opens the **Kafka Single Topic Partition Summary** display, which provides a view of the lag per partition, where the partitions are identified by the broker hosting the partition and the consumer reading the partition for a single topic, and also a list (table) of all topic partitions and their associated brokers and metrics.

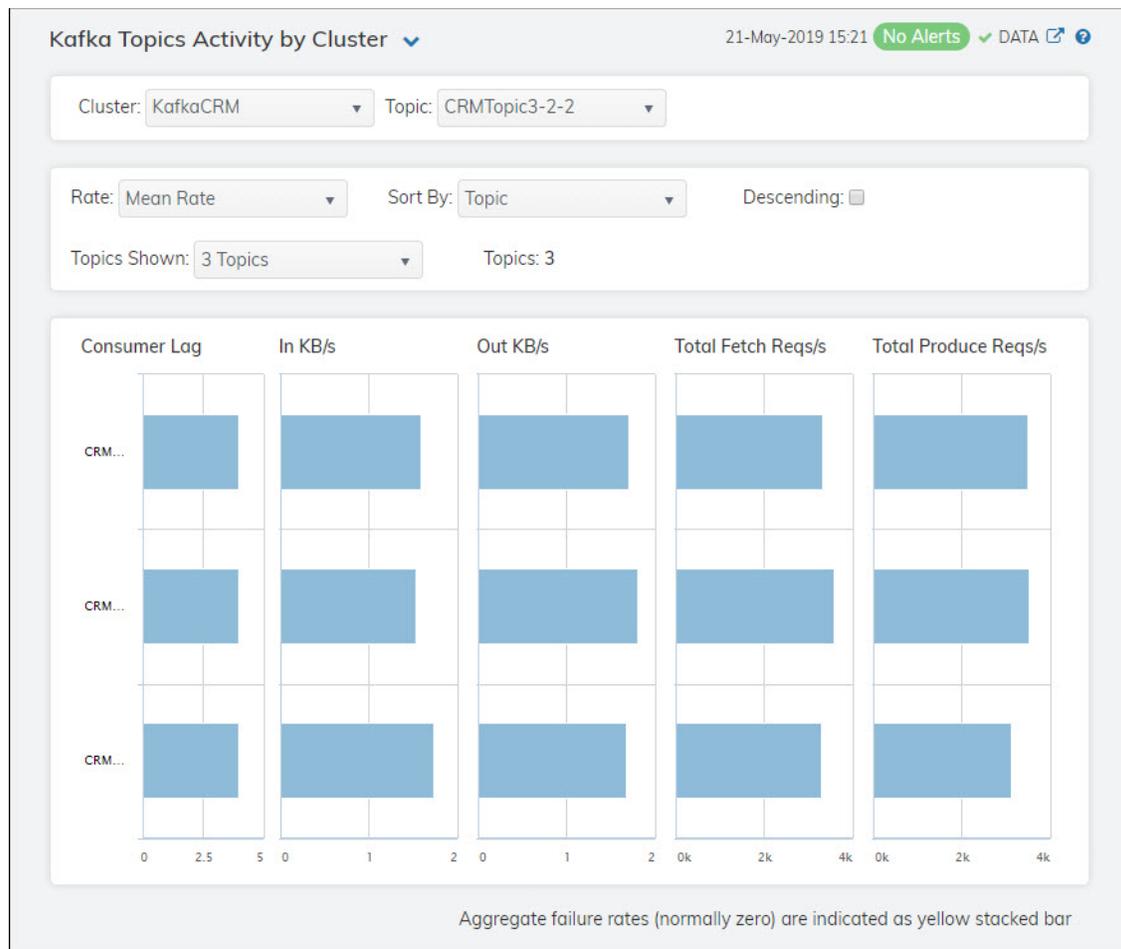


## Kafka Topics Activity by Cluster - HTML

Clicking **Topics Activity by Cluster** in the left/navigation menu opens the **Kafka Topics Activity by Cluster** display, which provides a view of the activity metrics on all topics for a particular cluster. You can view the metrics based on the mean rate, a 1 minute average rate, a 5 minute average rate, or a 15 minute average rate.

Select the metric from the **Sort By** drop down by which you want to sort the data in the display. When using this option with the **Sort Descending** toggle, the topics will be sorted in ascending or descending order using the option you select from this drop down. For example, if you select **In Msgs/s** from this drop down and select the **Sort Descending** toggle, the topics listed in the display will be sorted so that the topic with the most **In Msgs/s** will be listed at the top followed by the topic with the next most **In Msgs/s**, and so on.

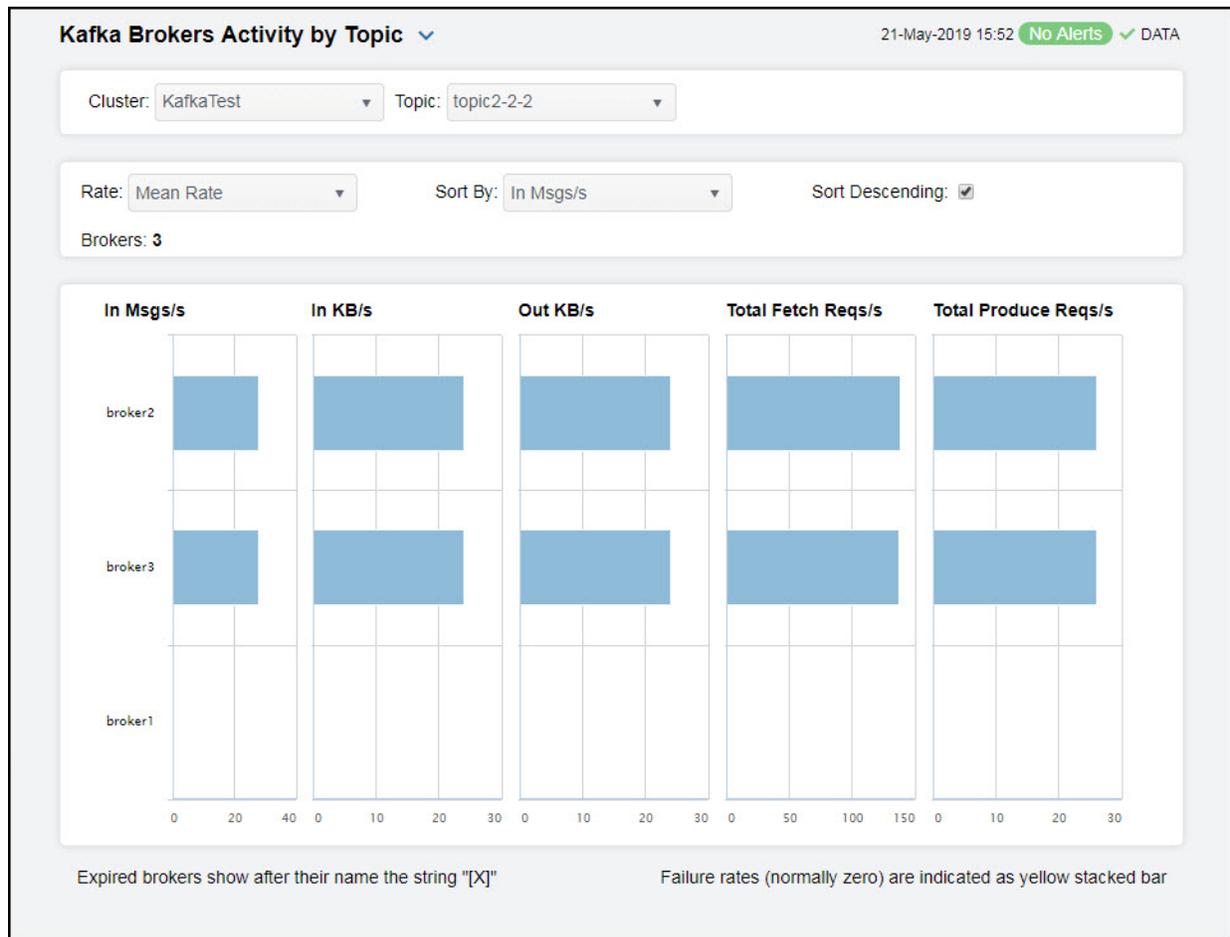
Select the option from the **Topics Shown** drop down to specify the number of topics you want to view.



## Kafka Brokers Activity by Topic - HTML

Clicking **Brokers Activity by Topic** in the left/navigation menu opens the **Kafka Brokers Activity by Topic** display, which provides a view of the activity metrics on all brokers for a particular topic. You can view the metrics based on the mean rate, a 1 minute average rate, a 5 minute average rate, or a 15 minute average rate.

Select the metric from the **Sort By** drop down by which you want to sort the data in the display. When using this option with the **Sort Descending** toggle, the brokers will be sorted in ascending or descending order using the option you select from this drop down. For example, if you select **In Msgs/s** from this drop down and select the **Sort Descending** toggle, the brokers listed in the display will be sorted so that the broker with the most **In Msgs/s** will be listed at the top followed by the broker with the next most **In Msgs/s**, and so on.



## Kafka Producers View - HTML

These displays provide detailed data for all producers or for a particular producer. Clicking **Kafka Producers** from the left/navigation menu opens the ["Kafka Producers Table - HTML"](#) display, which provides a view of all of your producers and their associated metric data. The options available under **Kafka Producers** are:

- **Producers Heatmap:** Opens the ["Kafka Producers Heatmap - HTML"](#) display, which allows you to view all producers and their associated metrics in a particular cluster.
- **Single Producer Summary:** Opens the ["Kafka Single Producer Summary - HTML"](#) display, which contains current and historical metrics, as well as trend data, for a single producer.
- **Single Producer JVM Runtime Summary:** Opens the ["Kafka Single Producer JVM Runtime Summary - HTML"](#) display, which contains current and historical JVM runtime metrics, as well as trend data, for a single producer.

## Kafka Producers Table - HTML

The table in this display provides a view of all of your producers and their associated metric data including connection, alert level, alert count, cluster name, client ID, and the current value of each gathered metric. Each row in the table contains data for a particular producer. Click a column header to sort column data in ascending or descending order. Double-click on a table row to drill-down to the ["Kafka Single Producer Summary - HTML"](#) display and view metrics for that particular producer. Toggle between the commonly accessed displays by clicking the drop down list on the display title.

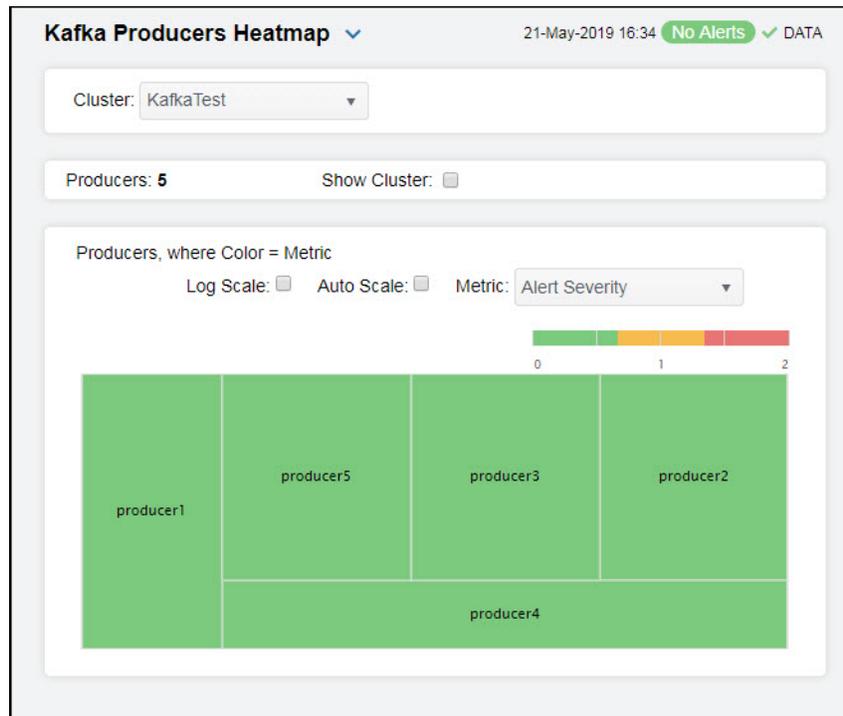
Cluster	Producer	Alert Level	Alert Count	Client ID
KafkaTest	producer1	✓		producer1
KafkaTest	producer5	✓		producer5
KafkaTest	producer3	✓		producer3
KafkaTest	producer2	✓		producer2
KafkaTest	producer4	✓		producer4

## Kafka Producers Heatmap - HTML

Clicking **Producers Heatmap** in the left/navigation menu opens the **Kafka Producers Heatmap**, which provides an easy-to-view interface that allows you to quickly identify the current status of each of your producers for each available metric. You can view the producers in the heatmap based on the following metrics: the current alert severity, the current alert count, the incoming/outgoing byte rate, the IO wait time, the request latency, and the request/response rates. By default, this display shows the heatmap based on the **Alert Severity** metric.

Each rectangle in the heatmap represents a producer. The rectangle color indicates the most critical alert state associated with the producer. Choose a cluster from the drop-down menu to view all producers for that cluster. Mouse over a rectangle to see additional metrics.

Drill-down and investigate a producer by clicking a rectangle in the heatmap to view details in the ["Kafka Single Producer Summary - HTML"](#) display.



### Available Metrics

Select the metric driving the heatmap display. The default is **Alert Severity**. Each **Metric** has a color gradient bar that maps values to colors. Mouse-over any rectangle to display the current values of the metrics for the producer. Click on a rectangle to drill-down to the associated ["Kafka Single Producer Summary - HTML"](#) display for a detailed view of metrics for that particular producer.

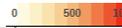
**Alert Severity** The current alert severity. Values range from **0** - **2**, as indicated in the color gradient  bar, where **2** is the highest Alert Severity:

-  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
-  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
-  Green indicates that no metrics have exceeded their alert thresholds.

**Alert Count** The total number of critical and warning unacknowledged alerts in the adapters. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.

**Incoming Bytes/s** The rate of incoming bytes (per second). The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **KafkaProducerIncomingByteRate**. The middle value in the gradient bar indicates the middle value of the range.

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

- Outgoing Bytes/s** The rate of outgoing bytes (per second). The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **KafkaProducerOutgoingByteRate**. The middle value in the gradient bar indicates the middle value of the range.
- When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.
- Avg IO Wait Time Avg ms** The average length of time the IO thread spent waiting for a socket (in milliseconds). The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **KafkaProducerIoWaitTimeMS**. The middle value in the gradient bar indicates the middle value of the range.
- When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.
- Avg Request Latency** The average amount of time between when a producer is called and when the producer receives a response from the broker. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **KafkaProducerRequestLatency**. The middle value in the gradient bar indicates the middle value of the range.
- When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.
- Requests/s** The average number of requests sent per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **KafkaProducerRequestRate**. The middle value in the gradient bar indicates the middle value of the range.
- When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.
- Responses/s** The average number of responses received (per second). The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **KafkaProducerResponseRate**. The middle value in the gradient bar indicates the middle value of the range.
- When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

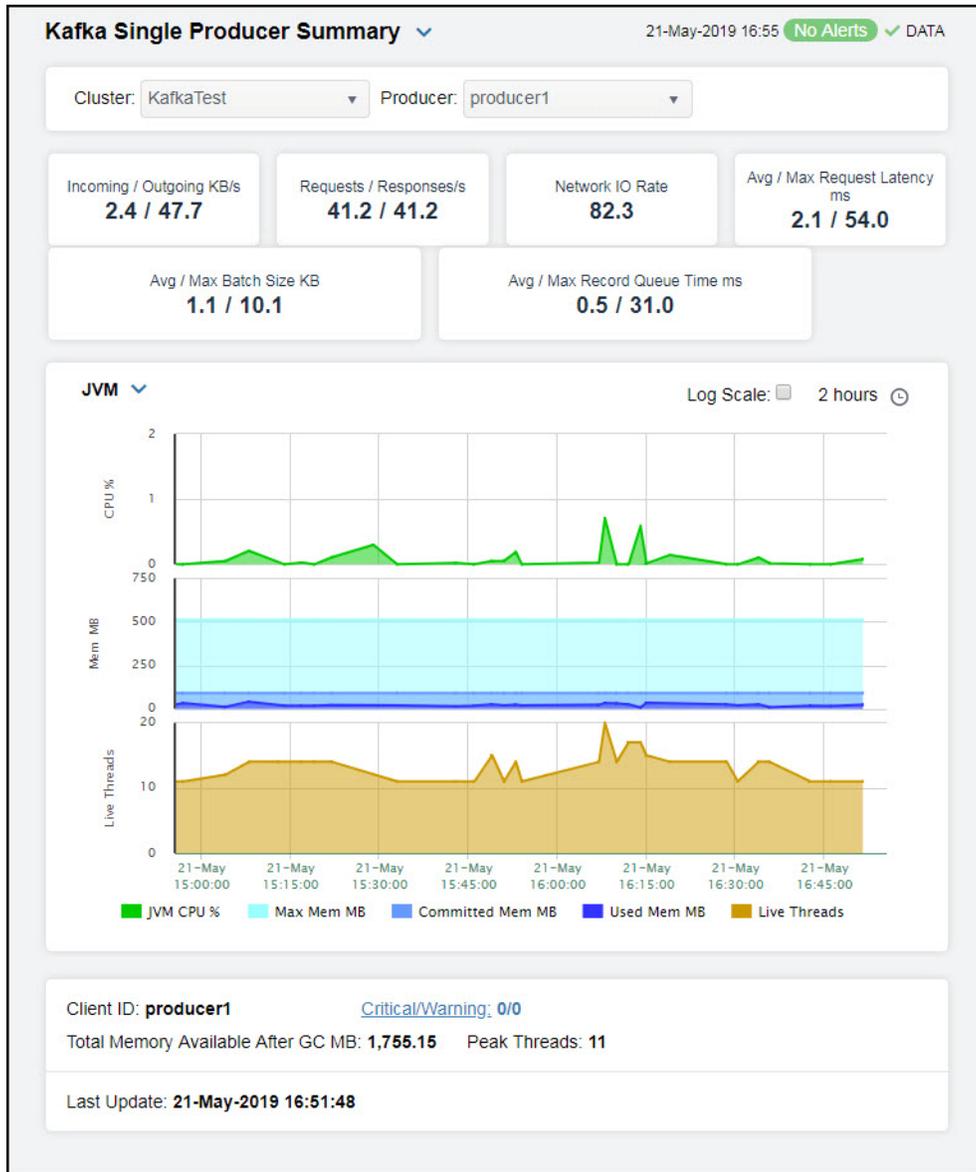
## Kafka Single Producer Summary - HTML

Clicking **Single Producer Summary** in the left/navigation menu opens the **Kafka Single Producer Summary** display, which provides a view of the current and historical metrics for a single producer.

Clicking on the information boxes at the top of the display takes you to the "[Kafka Producers Table - HTML](#)" display, where you can view additional producers data.

There are two options in the trend graph: **Performance** and **JVM**. In the **Performance** option on the trend graph, you can view trend data for the requests rate, responses rate, maximum and average request latency, outgoing kilobyte rate, and average IO wait time over a selected time range. In the **JVM** option on the trend graph, you can view trend data for JVM CPU percentage, maximum memory in megabytes, committed memory in megabytes, used memory in megabytes, and number of live threads over a selected time range.

Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display.



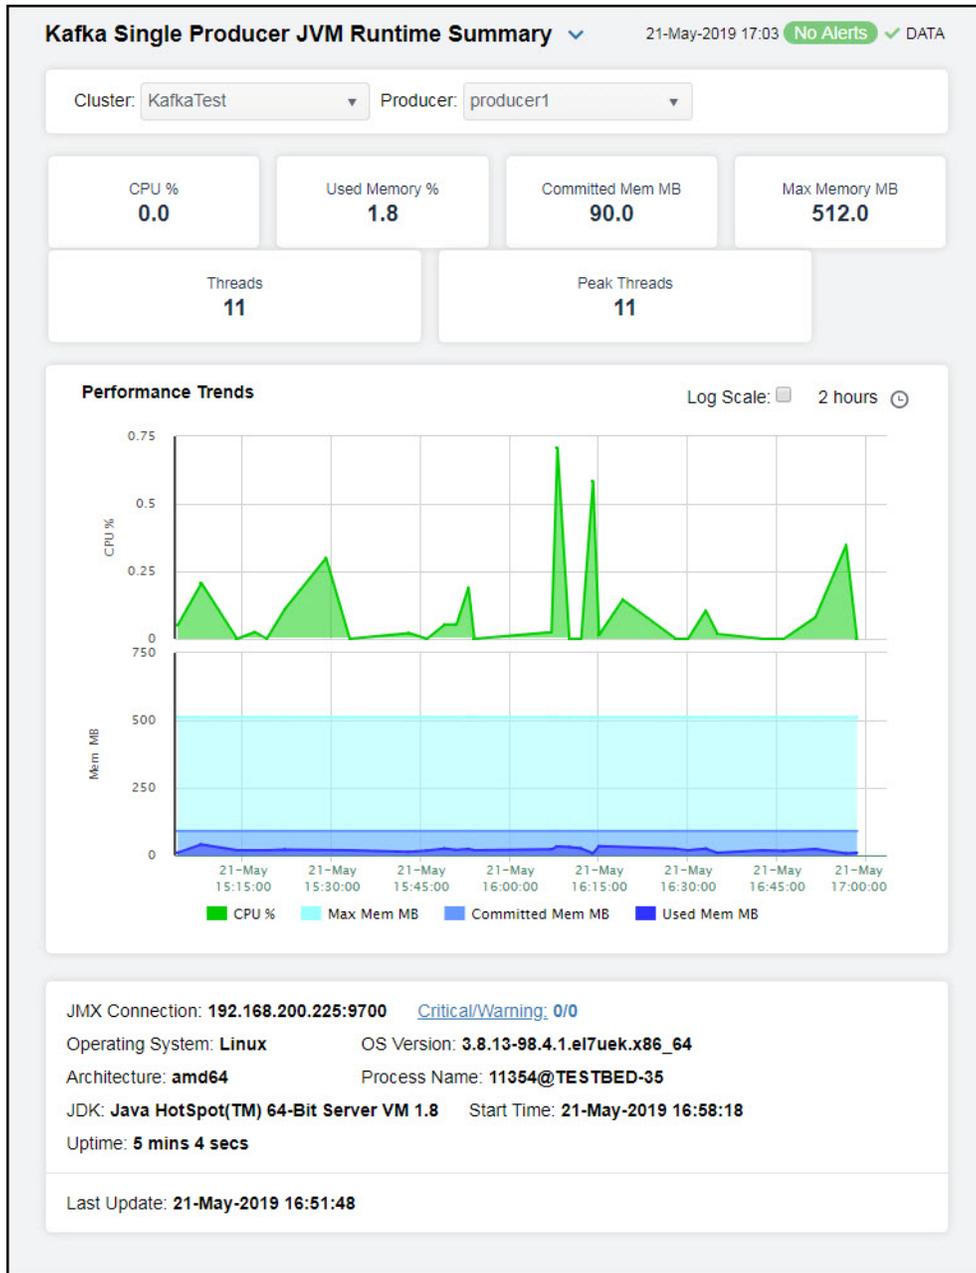
## Kafka Single Producer JVM Runtime Summary - HTML

Clicking **Single Producer JVM Runtime Summary** in the left/navigation menu opens the **Kafka Single Producer JVM Runtime Summary** display, which provides a view of JVM runtime statistics and trend data for the selected producer.

Clicking on the information boxes at the top of the display takes you to the "[Kafka Producers Table - HTML](#)" display, where you can view additional producers data.

The **Performance Trends** trend graph provides trend data for the CPU percentage, maximum memory in megabytes, committed memory in megabytes, and the used memory in megabytes over a selected time range.

Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display.



## Kafka Consumers View - HTML

These displays provide detailed data for all consumers or for a particular consumer. Clicking **Kafka Consumers** from the left/navigation menu opens the ["Kafka Consumers Table - HTML"](#) display, which provides a view of all of your consumers and their associated metric data. The options available under **Kafka Consumers** are:

- **Consumers Heatmap**: Opens the ["Kafka Consumers Heatmap - HTML"](#) display, which allows you to view all consumers and their associated metrics in a particular cluster.
- **Single Consumer Summary**: Opens the ["Kafka Single Consumer Summary - HTML"](#) display, which contains current and historical metrics, as well as trend data, for a single consumer.
- **Single Consumer JVM Runtime Summary**: Opens the ["Kafka Single Consumer JVM Runtime Summary - HTML"](#) display, which contains current and historical JVM runtime metrics, as well as trend data, for a single consumer.
- **Single Consumer Lag Summary**: Opens the ["Kafka Single Consumer Lag Summary - HTML"](#) display, which displays the lag per topic in a bar graph format and lists the lag per topic for the consumer.

## Kafka Consumers Table - HTML

The table in this display provides a view of all of your consumers and their associated metric data including connection, alert level, alert count, cluster name, client ID, and the current value of each gathered metric. Each row in the table contains data for a particular consumer. Click a column header to sort column data in ascending or descending order. Double-click on a table row to drill-down to the ["Kafka Single Consumer Summary - HTML"](#) display and view metrics for that particular consumer. Toggle between the commonly accessed displays by clicking the drop down list on the display title.

Kafka Consumers Table 22-May-2019 10:27 1 Alert ✓ DATA [↗](#) [?](#)

Cluster: - All -

Consumers: 6

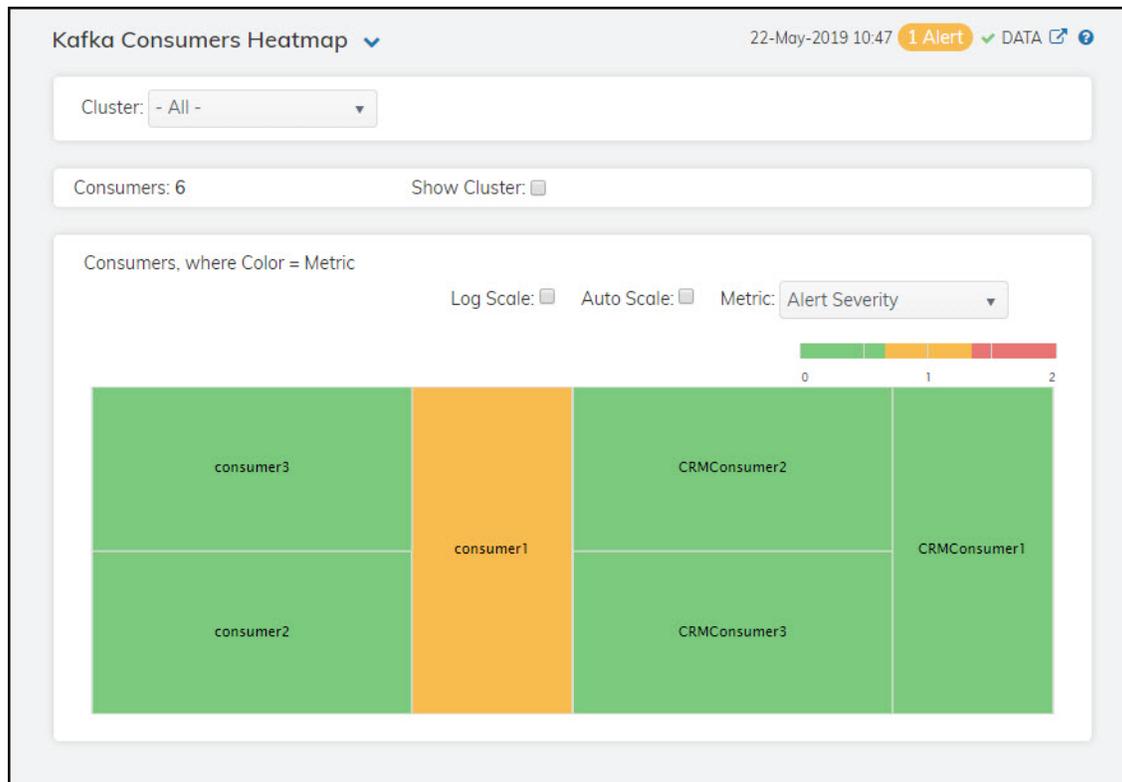
Cluster	Consumer	Alert Level	Alert Count	Client ID	Consumed Bytes/s
KafkaDemo	consumer3	✓	0	consumer3	0.0
KafkaCRM	CRMConsumer2	✓	0	CRMConsumer2	1,939.66
KafkaCRM	CRMConsumer3	✓	0	CRMConsumer3	1,930.85
KafkaCRM	CRMConsumer1	✓	0	CRMConsumer1	1,839.58
KafkaDemo	consumer2	✓	0	consumer2	0.0
KafkaDemo	consumer1	⚠	1	consumer1	3,080.29

## Kafka Consumers Heatmap - HTML

Clicking **Consumers Heatmap** in the left/navigation menu opens the **Kafka Consumers Heatmap**, which provides an easy-to-view interface that allows you to quickly identify the current status of each of your consumers for each available metric. You can view the consumers in the heatmap based on the following metrics: the current alert severity, the current alert count, the bytes consumed rate, the fetch latency average, the fetch rate, the maximum consumer lag, and the records consumed rate. By default, this display shows the heatmap based on the **Alert Severity** metric.

Each rectangle in the heatmap represents a consumer. The rectangle color indicates the most critical alert state associated with the consumer. Choose a cluster from the drop-down menu to view all consumers for that cluster. You can use the **Show Cluster** check-box  to include or exclude labels in the heatmap, and you can mouse over a rectangle to see additional metrics.

Drill-down and investigate a producer by clicking a rectangle in the heatmap to view details in the ["Kafka Single Consumer Summary - HTML"](#) display.



### Available Metrics

Select the metric driving the heatmap display. The default is **Alert Severity**. Each **Metric** has a color gradient bar that maps values to colors. Mouse-over any rectangle to display the current values of the metrics for the consumer. Click on a rectangle to drill-down to the associated ["Kafka Single Consumer Summary - HTML"](#) display for a detailed view of metrics for that particular consumer.

**Alert Severity** The current alert severity. Values range from **0 - 2**, as indicated in the color gradient bar, where **2** is the highest Alert Severity:

- Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
- Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
- Green indicates that no metrics have exceeded their alert thresholds.

**Alert Count** The total number of critical and warning unacknowledged alerts in the adapters. The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.

**Consumed Bytes/s** The rate of bytes being consumed (per second). The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **KafkaConsumer**. The middle value in the gradient bar indicates the middle value of the range.

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

- Fetch Latency Avg** The average time taken for fetch request. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **KafkaConsumer**. The middle value in the gradient bar indicates the middle value of the range.
- When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.
- Fetch Rate** The number of fetch request per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **KafkaConsumer**. The middle value in the gradient bar indicates the middle value of the range.
- When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.
- Consumer Lag** The maximum lag in the number of records for any partition. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **KafkaConsumer**. The middle value in the gradient bar indicates the middle value of the range.
- When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.
- Consumer Lag Rate** The lag rate in the number of records for any partition. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **KafkaConsumer**. The middle value in the gradient bar indicates the middle value of the range.
- When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.
- Consumed Records/s** The rate of records being consumed (per second). The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **KafkaConsumer**. The middle value in the gradient bar indicates the middle value of the range.
- When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

## Kafka Single Consumer Summary - HTML

Clicking **Single Consumer Summary** in the left/navigation menu opens the **Kafka Single Consumer Summary** display, which provides a view of the current and historical metrics for a single consumer.

Clicking on the information boxes at the top of the display takes you to the "[Kafka Consumers Table - HTML](#)" display, where you can view additional consumers data.

There are three options in the trend graph: **Performance**, **Consumption Rates**, and **JVM Resource Utilization**. In the **Performance** option on the trend graph, you can view trend data for the JVM CPU percentage, maximum record lag, consumer lag, and fetch rate over a selected time range. In the **Consumption Rates** option on the trend graph, you can view trend data for the consumed kilobyte rate and the consumed records rate over a selected time range. In the **JVM Resource Utilization** option on the trend graph, you can view trend data for JVM CPU percentage, maximum memory in megabytes, committed memory in megabytes, and used memory in megabytes over a selected time range.

Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display.



## Kafka Single Consumer JVM Runtime Summary - HTML

Clicking **Single Consumer JVM Runtime Summary** in the left/navigation menu opens the **Kafka Single Consumer JVM Runtime Summary** display, which provides a view of the current and historical JVM Runtime metrics for a single consumer.

Clicking on the information boxes at the top of the display takes you to the "[Kafka Consumers Table - HTML](#)" display, where you can view additional consumers data.

The **Performance Trends** trend graph shows trend data for the JVM CPU percentage, maximum memory in megabytes, committed memory in megabytes, and used memory in megabytes over a selected time range.

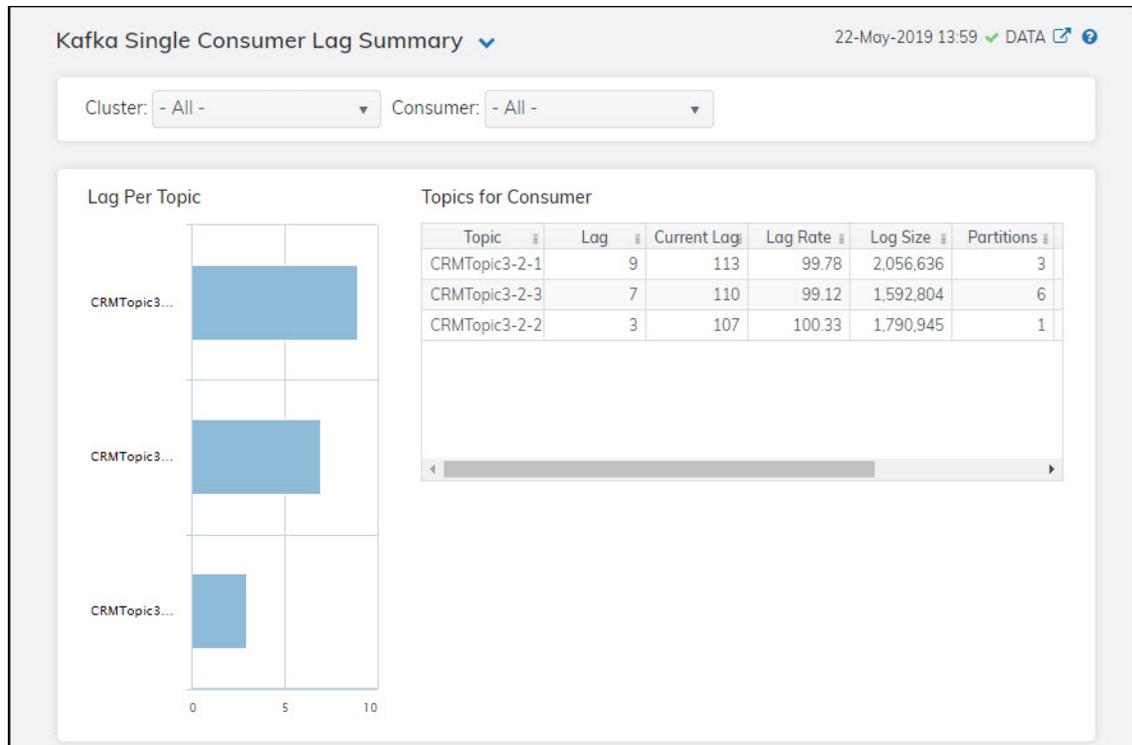
Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display.



## Kafka Single Consumer Lag Summary - HTML

Clicking **Single Consumer Lag Summary** in the left/navigation menu opens the **Kafka Single Consumer Lag Summary** display, which displays the lag per topic in a bar graph format and lists the lag per topic for the consumer. Double-click on a bar graph to drill-down to the ["Kafka Single Consumer Summary - HTML"](#) display and view metrics for that particular consumer.

Each row in the table contains data for a particular topic. Click a column header to sort column data in ascending or descending order. Toggle between the commonly accessed displays by clicking the drop down list on the display title.



## CHAPTER 6 RTView DataServer for Oracle

The RTView DataServer for Oracle provides a way to create connections and modify default configuration settings for solution packages and sends collected data to RTViewCentral. RTViewCentral contains the displays associated with the RTView DataServer for Oracle which you use to monitor your Oracle components. Both the display server user interface and the HTML user interface are described here.

The RTView *DataCollector* for Oracle is also available for use with the RTView DataServer for Oracle. RTView DataCollector for Oracle is used for collecting solution package data and sending it to one or more RTView DataServers. The RTView DataCollector for Oracle is useful if you need to distribute data collection.

For an overview and details about configuring RTView Enterprise, including RTViewCentral, RTView DataServers, RTView DataCollectors and solution packages, see the *RTView Enterprise Configuration Guide*.

RTViewCentral contains the following displays that will be populated with data collected via the RTView DataServer for Oracle:

- "Oracle Coherence"
- "Oracle Coherence - HTML"
- "Oracle Database"
- "Oracle Database - HTML"
- "Connector for Oracle Enterprise Manager"
- "Oracle WebLogic"
- "Oracle WebLogic - HTML"

**Note:** This document assumes familiarity with the products monitored. For additional details, refer to vendor documentation.

## Oracle Coherence

The Solution Package for Oracle Coherence provides information about the health and configuration of your Oracle Coherence cluster elements, including caches, nodes, services and clients.



The Solution Package for Oracle Coherence collects metrics from all your Coherence elements simultaneously, and does so at frequent intervals (typically every 10 seconds). At each interval, the OC Monitor performs analytic calculations on the gathered metrics (on the Data Server rather than a database for optimal performance) in terms of the cluster as a whole. It then presents consistently updated health "snapshots" of your entire cluster, in real time, using a dashboard format and visually rich and legible graphics.

The Solution Package for Oracle Coherence can be configured for a single Coherence cluster or multiple Coherence clusters. The Solution Package for Oracle Coherence is also often used in pre-production environments for conducting load testing and performance tuning.

The following Oracle Coherence Views can be found under **Components** tab > **Middleware** > **Oracle Coherence**:

- **"Cluster Selector"**: See all your Coherence clusters and Data Servers and choose which cluster to display data for.
- **"Cluster Views"**: Use these displays to assess Coherence cluster-level performance and utilization.
- **"Proxy Services"**: Use these displays to assess proxy service performance metrics.
- **"Cache Services"**: Use these displays to assess performance and utilization of all caches in the cluster.
- **"Federated Clusters"**: See all your Federated clusters and Data Servers and choose which cluster to display data for.

- **"All Caches"**: Use these displays to investigate performance, utilization and activity metrics of a single cache.
- **"Single Cache"**: Use these displays to assess node-level performance and utilization in the cluster.
- **"All Nodes"**: Use these displays to investigate performance and utilization metrics of a single node.
- **"Single Node"**: Use these displays to investigate performance and utilization metrics of a single node.
- **"Time Range Analysis"**: Use these displays to manage your Oracle Coherence metrics, nodes and caches.
- **"OC Administration"**: Use these displays to manage your Oracle Coherence metrics, nodes and caches.

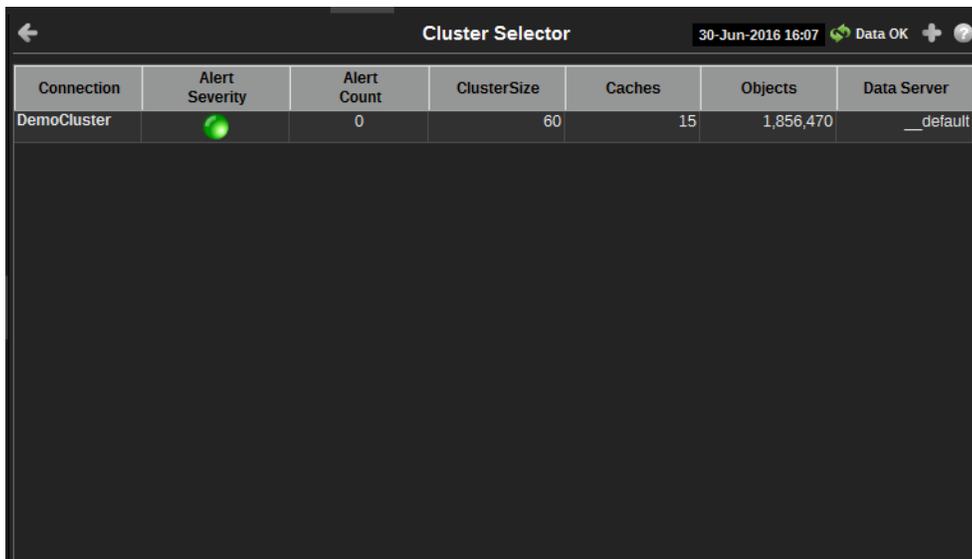
## Cluster Selector

This display shows details about your Coherence clusters and Monitor Data Servers.

Use this display to see all the Coherence clusters you can monitor, as well as their status. Choose a cluster to view performance details for the cluster in the **"Cluster - Overview"** display.

Each row in the table is a different Coherence cluster. The columns contain information pertaining to each cluster. When you select a cluster you are also selecting the Data Server corresponding with that cluster. After you make your selection, all displays subsequently show data for that cluster/Data Server (except for alert displays which consolidate alerts from all Data Servers). For example, the **"Node Summary"** display will then show data for the selected cluster/Data Server.

For details about Oracle Coherence data, refer to vendor documentation at [www.oracle.com](http://www.oracle.com).



The screenshot shows a mobile application interface titled "Cluster Selector". At the top right, it displays the date and time "30-Jun-2016 16:07" and a status indicator "Data OK" with a green checkmark. Below the title is a table with the following columns: Connection, Alert Severity, Alert Count, ClusterSize, Caches, Objects, and Data Server. The table contains one row for "DemoCluster". The "Alert Severity" column for "DemoCluster" has a green circle icon. The "Data Server" column for "DemoCluster" has the value "\_\_\_default".

Connection	Alert Severity	Alert Count	ClusterSize	Caches	Objects	Data Server
DemoCluster		0	60	15	1,856,470	___default



<b>Connection</b>	The name of the user defined connection that is used to connect to the monitored Coherence cluster.
<b>Alert Severity</b>	The maximum level of alerts on the cluster. <ul style="list-style-type: none"> <li>Red indicates that one or more exceeded their ALARM LEVEL threshold.</li> <li>Yellow indicates that one or more exceeded their WARNING LEVEL threshold.</li> <li>Green indicates that none have exceeded their alert thresholds.</li> </ul>
<b>Alert Count</b>	The number of current alerts for the cluster.
<b>Cluster Size</b>	The total number of nodes for the cluster.
<b>Caches</b>	The total number of caches for the cluster.
<b>Objects</b>	The total number of objects stored in the cluster.
<b>Data Sever</b>	The name of the Data Server (connection) that is used to monitor the cluster.

## Cluster Views

Cluster Views displays present high-level performance metrics for the cluster. Use the Cluster Views displays to quickly assess Coherence cluster-level performance metrics.

- ["Cluster - Overview" on page 695](#): Quickly assess general cluster stability, cluster size (number of nodes, clients and caches), service and cache capacity utilization/distribution and HA status.
- ["Caches / Nodes / Alerts" on page 699](#): View cache and node utilization hot spots and currently active alerts.
- ["Memory/Network Health" on page 702](#): Assess cluster memory utilization and packet transmission success/failure trends, and see weakest nodes.
- ["Stability Metrics" on page 704](#): Troubleshoot nodes joining and leaving the cluster, view HA status for cache services.
- ["All Services History" on page 706](#): Assess capacity utilization, over time, by all services in a cluster.
- ["All Caches History" on page 709](#): Assess capacity utilization and distribution for all caches in a cluster, and quickly identify potential bottlenecks.
- ["All Nodes History" on page 713](#): Assess capacity utilization, over time, for all nodes in a cluster.

## Cluster - Overview

Use this display to quickly assess the cluster size (number of nodes, clients and caches) and stability, service and cache capacity utilization and HA status. This display is the initial view in the Monitor.

Choose a cluster from the drop down menu. Check the Communication Success% bar charts for cluster packet loss. If the pairs of bar graphs are uneven, this indicates that packet loss is occurring. The cause for the packet loss could be a network issue, a single defective NIC card, a garbage collection issue, disk swapping or a shortage of CPU on a single machine. Investigate further by clicking the bar chart to view details in the Cluster - "Memory/Network Health" display.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

### Fields and Data:

#### Coherence Cluster Configuration

- Total Nodes** Total number of nodes being monitored, including storage enabled nodes, client nodes, and management (JMX) nodes.
- Storage** Total number of nodes in the cluster which have storage enabled for any cache. This value is equal to the total nodes when replicated caches are being used. The number is less when only distributed cache types are utilized.

<b>Clients</b>	Total number of nodes in the cluster which do not have storage enabled for any cache. These are usually process nodes, proxy nodes, extend nodes, or MBean server nodes.
<b>Caches</b>	Total number of caches in the cluster.
<b>Version</b>	Version of Oracle Coherence running.
<b>Cluster Memory Usage Totals</b>	
<b>Senior Node</b>	Node ID of the senior node of the cluster.
<b>Client Nodes</b>	Monitor client node memory utilization for the cluster.
	<b>Max MB</b> Total memory allocated.
	<b>Used MB</b> Total memory used.
	<b>%</b> <b>Percent of allocated memory being used.</b>
<b>Storage Nodes</b>	Monitor storage node memory utilization for the cluster.
	<b>Max MB</b> Total memory allocated.
	<b>Used MB</b> Total memory used.
	<b>%</b> <b>Percent of allocated memory being used</b>
<b>Alert Severity</b>	The maximum level of alerts for all nodes in the cluster. Click to drill down to the Alert Detail Table.
	 Red indicates that one or more exceeded their ALARM LEVEL threshold.
	 Yellow indicates that one or more exceeded their WARNING LEVEL threshold.
	 Green indicates that none have exceeded their alert thresholds.
	<b>Memory</b> Represents the current most critical state of alerts related to heap and memory alerts for all nodes in the cluster. For example, the AvailableMemoryLowNode alert.
	<b>Network</b> Represents the current most critical state of alerts related to network and communication protocols for all nodes in the cluster. For example, the BadCommunicationCluster alert.
	<b>Stability</b> Represents the current most critical state of alerts related to cluster stability for all nodes in the cluster. For example, the DepartedNodePercentage alert.
	<b>Tasks</b> Represents the current most critical state of alerts related to queries, entry processors and invocations for all nodes in the cluster. For example, the HighTaskBacklogNode alert.
	<b>Data Quality</b> Represents the current most critical state of alerts related to the quality of data in the Data Server for all nodes in the cluster. For example, the JmxProcessingTime alert.
	<b>Other</b> Represents the current most critical state of alerts related to all alerts not represented in the other five status indicators for all nodes in the cluster. For example, the CapacityLimitAllCaches alert.
	<b>Memory</b> Represents the current most critical state of alerts related to heap and memory alerts for all nodes in the cluster. For example, the AvailableMemoryLowNode alert.

**Service Configuration & HA Status**

<b>Cache Services</b>	Assess size, distribution and status of Coherence protocol-related cache services used by applications in the cluster. Determine whether cache services are distributed properly across the cluster. The list includes distributed, replicated and mirrored caches. Note that Management and Invocation services are intentionally not listed.
<b>Service Name</b>	The name of the service in the cluster. These are defined in each server cache configuration XML file.
<b>StatusHA</b>	The high availability status for each of the services.
<b>MACHINE-SAFE</b>	If a machine for the service goes offline the data stored on the machine remains available in the cluster (no data loss).
<b>NODE-SAFE</b>	If a node for the service goes offline (or is taken offline using kill-9) data stored on the node remains available in the cluster (no data loss).
<b>ENDANGE RED</b>	If a node for the service goes offline the data stored on the node is potentially unavailable in the cluster (potential data loss).
<b>Total Nodes</b>	The number of nodes in the cluster that are running a thread for the service.
<b>Storage Nodes</b>	The number of nodes for the service where storage is enabled.
<b>Caches</b>	The number of caches for the service.
<b>Objects</b>	The number of objects in all caches for the service.
<b>Senior</b>	The node ID of the most senior node in the cluster for the service.

### Caches - Busiest & Largest

<b>Most Gets</b>	Track services performing the greatest number of gets in the cluster. The total is the number of gets by nodes in the cluster since the last sample was retrieved. Click to drill-down to the All Caches - " <a href="#">Current Activity Chart</a> " display.
<b>Cumulative</b>	Select the checkbox to show only the cumulative total for all nodes for the service since they started in the Most Gets bar chart.
<b>Largest Cache</b>	Track caches that consume the greatest amount of capacity. Click to drill-down to the All Caches - " <a href="#">Current Size Chart</a> " display.

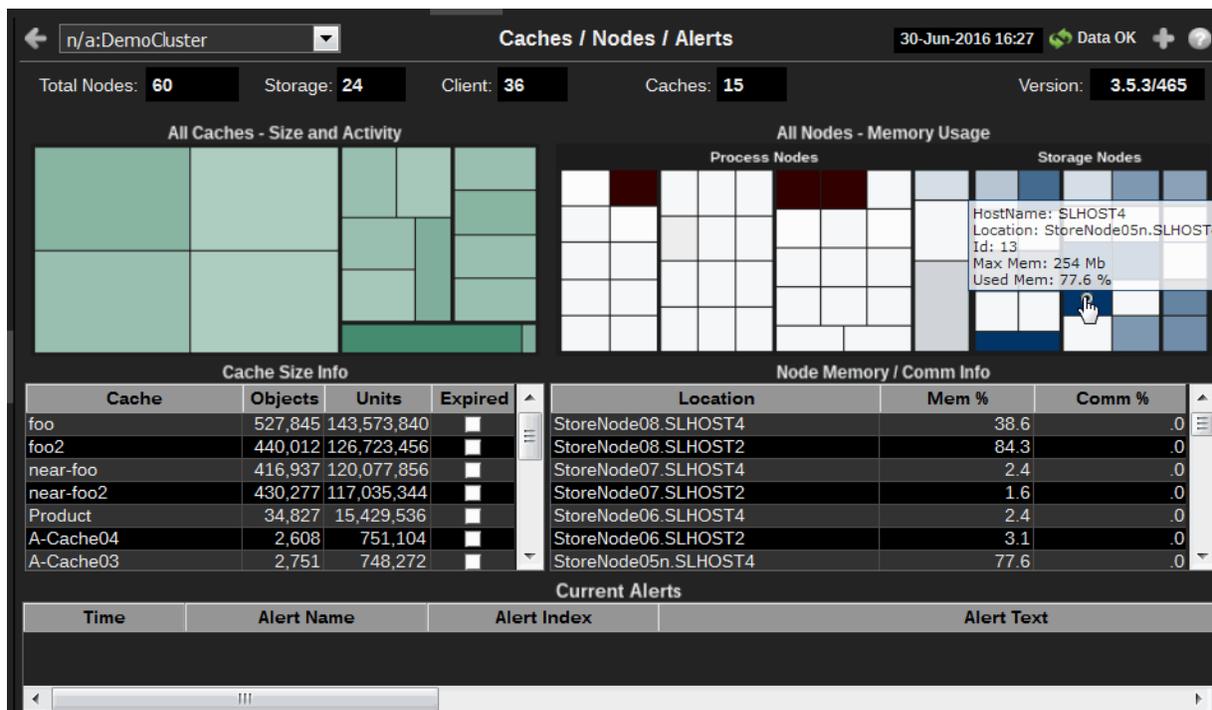
### Cluster Stability

<b>Node Uptimes</b>	Monitor cluster stability and how often nodes are restarted (for example, every month, every day, every hour, and so forth). If the number of nodes running for seconds of time increases (and your nodes are restarted weekly), consider investigating. Click in the Node Uptimes region to view details on the " <a href="#">Stability Metrics</a> " display.  Solid colors in the graph indicate the amount of time since the nodes were started. Longer uptimes generally represent a more stable cluster. Departed Nodes specifies the number of nodes that have departed and not returned since monitoring of the cluster was started. If a node departs and returns with the same name, the count is decremented.
<b>Memory Utilization%</b>	Monitor memory utilization for all nodes in the cluster.
<b>Average</b>	The average memory utilization for all nodes in the cluster.

	<b>Worst Node</b>	The most amount of memory consumed by a single node in the cluster. A slow node that provides data to other nodes can cause latency issues for the entire cluster. If a node is consuming too much memory, investigate by clicking the bar chart to view details in the Cluster - " <a href="#">Memory/Network Health</a> " display.
<b>Communication Success%</b>	<p>Monitor cluster packet loss--an excellent indicator of systemic issues in the cluster. If the pairs of bar graphs are uneven, this indicates that packet loss is occurring and analysis is needed. Investigate further by clicking the bar chart to view details in the Cluster - "<a href="#">Memory/Network Health</a>" display.</p> <p>The bar charts show the percent (%) successful UDP packet transfers in the cluster for the last twenty minutes. Each pair of bars show the Publish and Receive success rates for all nodes in the cluster. Compare each pair of Publish and Receive bars. The bars should have similar rates. If they do not have similar rates this indicates packet loss in the cluster. For example, if the Publish success rate is much lower than the Receive success rate, packets are being resent and the receiver is not getting them.</p> <p>Compare and track the pairs of bars across twenty minutes. The bars should track evenly. If the bars do not track evenly this also is a sign of packet loss in the cluster.</p> <p>The cause for the packet loss could be a network issue, a single defective NIC card, a garbage collection issue, disk swapping or a shortage of CPU on a single machine.</p>	<p><b>Publish</b> The Publish success rate is the percent (%) of packets in the cluster successfully sent by nodes, without having to be resent. A 100% success rate occurs when a packet is sent and does not have to be re-sent. When a packet must be resent the success rate is reduced.</p> <p><b>Receive</b> The Receive success rate is the percent (%) of packets in the cluster successfully received by nodes, without being received twice. A 100% success rate occurs when a packet is received once. When a packet is received twice the success rate is reduced.</p>

## Caches / Nodes / Alerts

Use this display to view cache and node utilization hot spots and currently active alerts. Observe how much capacity is taken from memory and how much is taken from consumption. Identify caches and nodes that are slow due to a shortage of capacity or memory. Verify nodes are configured properly (using the mouseover tool-tip). View time-ordered list of current alerts in the cluster.



**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

### Fields and Data:

- Total Nodes** Total number of nodes being monitored, including storage enabled nodes, client nodes, and management (JMX) nodes.
- Storage** Total number of nodes in the cluster which have storage enabled for any cache. This value is equal to the total nodes when replicated caches are being used. The number is less when only distributed cache types are utilized.
- Clients** Total number of nodes in the cluster which do not have storage enabled for any cache. These are usually process nodes, proxy nodes, extend nodes, or MBean server nodes.
- Caches** Total number of caches in the cluster.
- Version** Version of Oracle Coherence running.

## Capacity & Memory Usage

### All Caches - Size and Activity

Use the heatmap to identify a cache with high capacity or memory usage, indicated by a dark rectangle. Observe how much capacity is taken from memory and how much is taken from consumption. View cache metrics using the mouseover tool-tip. Investigate cache utilization trends over time in the ["All Caches History" display](#). Click on a rectangle to drill-down to the All Caches - ["All Caches Heatmap"](#).

The heatmap is grouped by service. Each rectangle represents a cache within the service. The size of each rectangle represents the size of a cache in units. The color of each rectangle represents the number of gets on the cache. The color is linearly scaled, where white is the minimum gets seen and dark green is the maximum gets seen.

#### Cache Size Info

The table lists each cache in the cluster and enables you to sort the by most/least amount of objects or units. Click a row to view details in the ["Single Cache Summary" display](#).

#### Cache

The name of the cache.

#### Objects

The number of objects currently in the cache.

#### Units

The number of units currently used by the cache.

### All Nodes- Memory Usage

Use the heatmap to identify a node with high memory usage, indicated by a dark rectangle. Verify nodes are configured properly using the mouseover tool-tip. Click on a rectangle to drill-down to the ["All Nodes by Type/Host/Memory"](#).

The heatmap is divided into two sections: Process Nodes and Storage Nodes. Each rectangle represents a node in the cluster. The size of the rectangle represents the value of the maximum node memory. The color of the rectangle represents the value of the memory used. The color is linearly scaled, where white is 0% memory used and dark green is 80% memory used.

#### Node Memory/Comm Info

The table lists each node in the cluster and enables you to sort the by most/least amount of objects or units. Click a row to view details in the ["Node Summary" display](#).

**Location** A unique identifier for each node. It is defined as: **member\_name.machine.rack.site**.

**Mem%** The percent memory utilization for the node.

**Comm%** The percent memory utilization used for packet transfer by the node.

## All Active Alerts (in selected cluster)

### Current Alerts

The table lists all alerts for all sources (nodes and caches) in the selected cluster that have exceeded an alert threshold. Sort the data by column using the button. By default, critical and warning alerts are shown. Select an alert in the list to open the **Alert Detail Table** dialog and acknowledge an alert or add comments. Where:

● Red indicates that one or more resources exceeded their ALARM LEVEL threshold.

● Yellow indicates that one or more resources exceeded their WARNING LEVEL threshold.

● Green indicates that no resources have exceeded their alert thresholds.

For details about alerts, see **Appendix, Alert Definitions**.

#### Alert Name

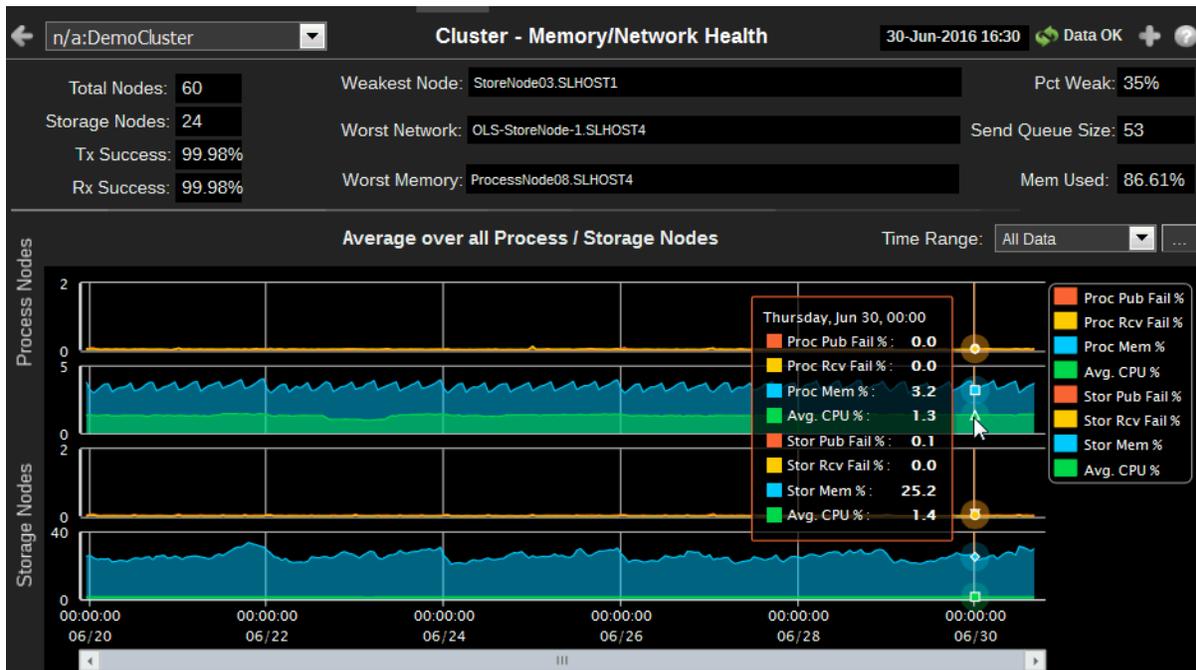
The alert type. Alert Types contain alert threshold definitions. A single alert type applies to all nodes or caches in the cluster. For example, the `OcAvailableMemoryLowNodeSpike` alert type applies to multiple nodes, and the `OcCapacityLimitCache` alert type applies to multiple caches. (The Alert Index identifies the source node for the alert.)

For details about alerts, see **Appendix, Alert Definitions**.

<b>Alert Index</b>	The Oracle Coherence source (node or cache) from which the alert originated. As with nodes, a cluster can have multiple caches. A single alert type, such as OcCapacityLimitCache, applies to all caches in the cluster. The Alert Index identifies the cache from which the alert originated.
<b>Alert Text</b>	Descriptive information about the alert.
<b>Cleared</b>	The checkbox is selected if this alert has cleared. An alert is considered cleared when the source for the alert (node or cache) returns to below the alert threshold. To include acknowledged alerts in the table, select Show Cleared.
<b>Acknowledged</b>	The checkbox is selected if this alert has been acknowledged. Acknowledged alerts have been manually acknowledged by an administrator. Acknowledged alerts are automatically removed from the Current Alerts table. To include acknowledged alerts in the table, select Show Acknowledged.
<b>ID</b>	Unique ID for the alert.
<b>Comments</b>	Comments about the alert previously entered by an administrator.
<b>Cleared Reason</b>	An alert is in a cleared state when the source for the alert (node or cache) returns to below the alert threshold. Or, with the OcDepartedNode alert type, when the node rejoins the cluster the alert is cleared.
<b>Cleared Time</b>	The time the alert was cleared.
<b>Alert Index Value</b>	The Oracle Coherence source (node or cache) from which the alert originated.
<b>Cluster Connection</b>	The name of the cluster in which the alert source (node or cache) is a member.

## Memory/Network Health

Use this display to assess cluster memory utilization and packet transmission success/failure trends, and to see the weakest nodes.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

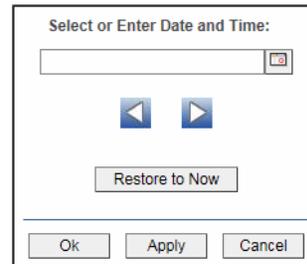
23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

### Fields and Data:

- Total Nodes** The total number of nodes in the cluster. This includes storage enabled nodes, client nodes, and management (JMX) nodes.
- Storage Nodes** The total number of nodes in the cluster which have storage enabled for any cache. This value is equal to the total nodes when replicated caches are being used. The number is less when only distributed cache types are utilized.
- Tx Success** The publisher success rate, in percent. The Publish success rate is the percent (%) of packets in the cluster successfully sent by nodes, without having to be resent. A 100% success rate occurs when a packet is sent and does not have to be re-sent. When a packet must be resent the success rate is reduced.
- Rx Success** The receiver success rate, in percent. The Receive success rate is the percent (%) of packets in the cluster successfully received by nodes, without being received twice. A 100% success rate occurs when a packet is received once. When a packet is received twice the success rate is reduced.
- Weakest Node** The node voted by Coherence as the weakest in the cluster. The Weakest Node often points to a server/node that is causing performance issues. The node value most often appears in the "weakest node" attribute of all the JMX "node" objects. The format of this string is **<Node IP Address>:< Node Port >/<NodeID>**.

	<b>Weak</b>	The percent of the Coherence nodes that "elected" the node as the weakest.
<b>Worst Network</b>		The node that has the longest network queue in the cluster.
	<b>Send Queue</b>	The number of packets currently scheduled for delivery, including packets sent and still awaiting acknowledgment. Packets that do not receive an acknowledgment within the ResendDelay interval are automatically resent.
<b>Worst Memory</b>		The node that has the lowest available memory of any node in the cluster.
	<b>Mem Used</b>	The percent of memory consumed on the Worst Memory node.
<b>Average over all Process / Storage Nodes</b>	<b>Trend Graphs</b>	The trend graphs show aggregated performance metrics for storage and process nodes.

**Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

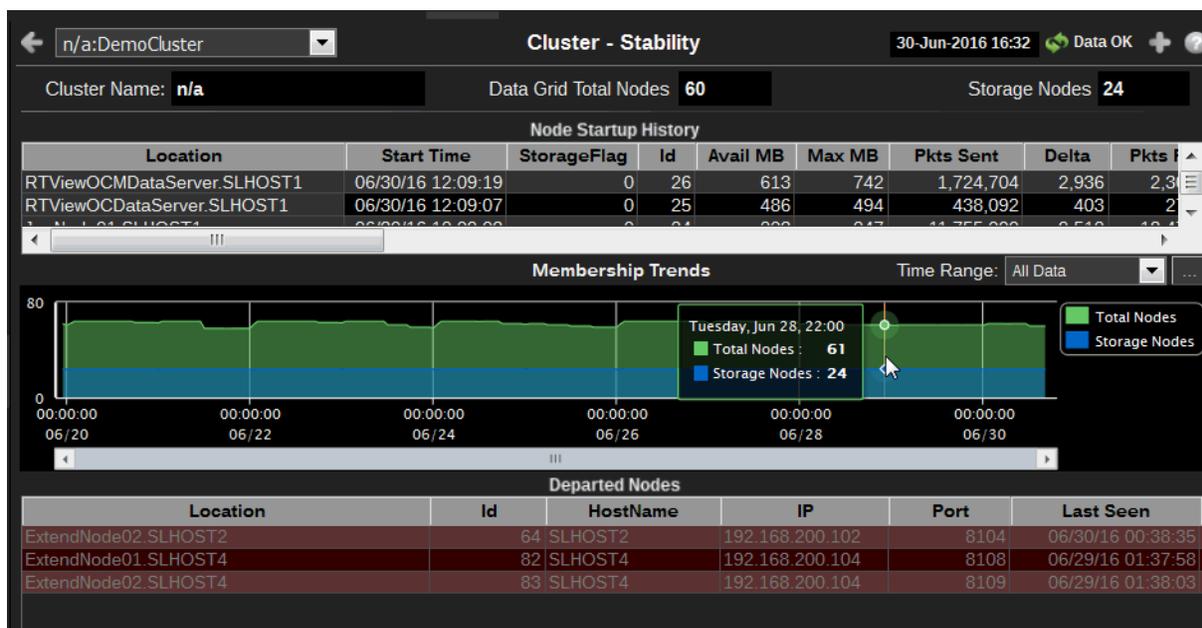
Click **Restore to Now** to reset the time range end point to the current time.

<b>Process Nodes</b>	Publish Failures and Received Failures	Indicates the trending of process node publisher and receiver failure rates. If these values are above 10%, action may be required to improve the stability or performance of the cluster as a whole. The Weakest Node information often points to the server/nodes that are the cause of these issues.
	Memory Utilization%	Indicates the trending of process node memory utilization. If these values are above 10%, action may be required to improve the stability or performance of the cluster as a whole.

<b>Storage Nodes</b>	Publish Failures and Received Failures	Indicates the trending of storage node publisher and receiver failure rates. If these values are above 10%, action may be required to improve the stability or performance of the cluster as a whole. The Weakest Node information often points to the server/nodes that are the cause of these issues.
	Memory Utilization%	Indicates the trending of storage node memory utilization. If these values are above 10%, action may be required to improve the stability or performance of the cluster as a whole.

## Stability Metrics

Use this display to troubleshoot nodes joining and leaving the cluster, and view HA status for cache services. This display presents information about node up times and the stability of the cluster.



**Title Bar (possible features are):**

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu Table open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

### Fields and Data:

<b>Cluster Name</b>	Select a cluster from the drop-down menu.
<b>Data Grid Total Nodes</b>	The total number of nodes being monitored. This includes storage enabled nodes, client nodes, and management (JMX) nodes.

**Storage Nodes**

The total number of nodes in the cluster which have storage enabled for any cache. This value is equal to the total nodes when replicated caches are being used. The number is less when only distributed cache types are utilized.

**Node Startup History**

Use this table to identify nodes that have departed and returned to the cluster recently. This table contains a list of nodes in the cluster, sorted by start time (the most recently created node is listed first).

<b>Location</b>	A unique identifier for each node. It is defined as: <b>member_name.machine.rack.site</b> .
<b>Start Time</b>	The date and time that the node joined the cluster.
<b>StorageFlag</b>	Indicates whether storage is enabled ( <b>0</b> or <b>1</b> ).
<b>Id</b>	The short member id that uniquely identifies this member.
<b>Avail MB</b>	The amount of available memory for this node, in megabytes.
<b>Max MB</b>	The maximum amount of memory for this node, in megabytes.
<b>Pkts Sent</b>	The cumulative number of packets sent by this node since the node statistics were last reset.
<b>Delta</b>	The number of packets sent by this node since the last update.
<b>Pkts Rcvd</b>	The cumulative number of packets received by this node since the node statistics were last reset.
<b>Delta</b>	The number of packets received by this node since the last update.
<b>Pkts Rptd</b>	The cumulative number of duplicate packets received by this node since the node statistics were last reset.
<b>Delta</b>	The number of duplicate packets received by this node since the last update.
<b>Pkts Resent</b>	The cumulative number of packets resent by this node since the node statistics were last reset.
<b>Delta</b>	The number of packets resent by this node since the last update.
<b>Pub Succ Rate</b>	The publisher success rate for this node since the node statistics were last reset. Publisher success rate is a ratio of the number of packets successfully delivered in a first attempt to the total number of sent packets. A failure count is incremented when there is no ACK received within a timeout period. It could be caused by either very high network latency or a high packet drop rate.
<b>Rec Succ Rate</b>	The receiver success rate for this node since the node statistics were last reset. Receiver success rate is a ratio of the number of packets successfully acknowledged in a first attempt to the total number of received packets. A failure count is incremented when a re-delivery of previously received packet is detected. It could be caused by either very high inbound network latency or lost ACK packets.
<b>Member</b>	The member name for this node.
<b>Machine</b>	The machine name for this node.
<b>Rack</b>	The rack name for this node.
<b>Site</b>	The site name for this node.
<b>Process</b>	The process name for this node.
<b>Uni Addr</b>	The unicast address. This is the IP address of the node's DatagramSocket for point-to-point communication.

<b>Uni Port</b>	The unicast port. This is the port of the node's DatagramSocket for point-to-point communication.
<b>RoleName</b>	The role name for this node.
<b>Product-Edition</b>	The product edition this node is running. Possible values are: Standard Edition (SE), Enterprise Edition (EE), Grid Edition (GE).
<b>Membership Trends</b>	Track the total number of nodes and the total number of storage nodes in the cluster for the duration of the user session. These lines are normally unchanging or "flat". If there are fluctuations in this graph, check the debugging guide for appropriate actions.
<b>Time Range</b>	Select a time range from the drop down menu varying from <b>2 Minutes</b> to <b>Last 7 Days</b> , or display <b>All Data</b> . To specify a time range, click Calendar  .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

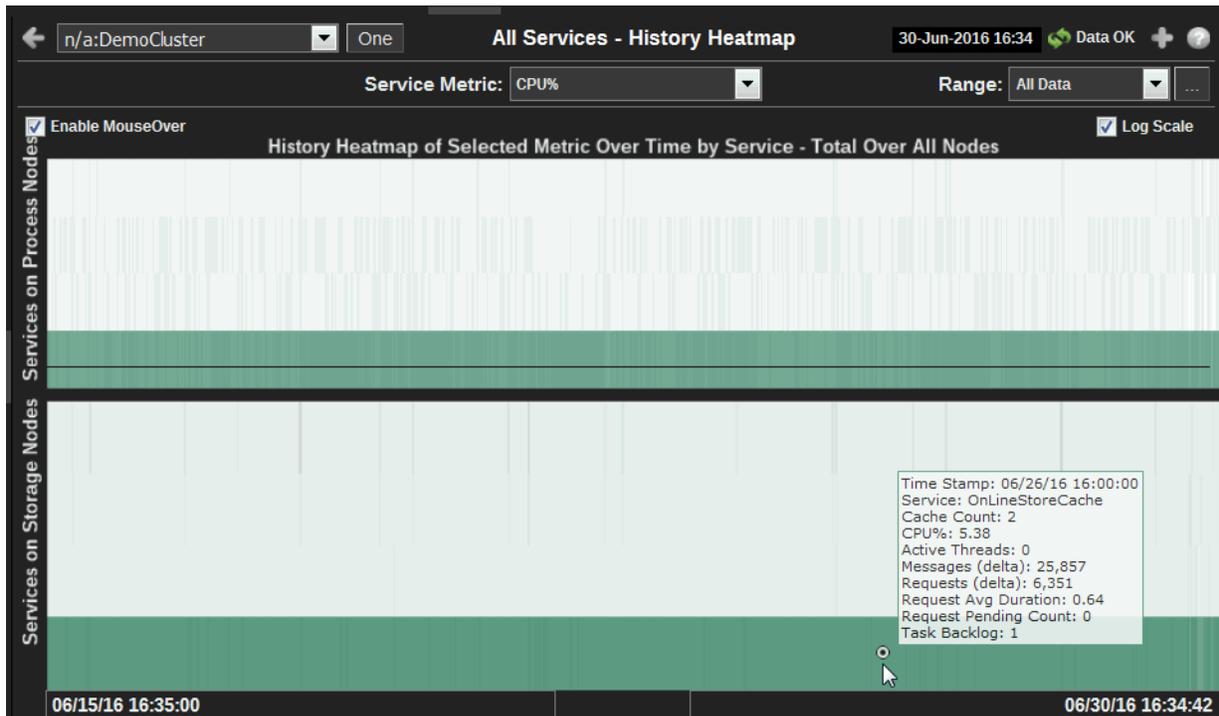
<b>Departed Nodes</b>	Track departed nodes by IP address, port number and time last seen.
<b>Location</b>	A unique identifier for each node. It is defined as: <b>member_name.machine.rack.site</b> .
<b>HostName</b>	The name of the host on which the node resides.
<b>IP</b>	The node IP address.
<b>Port</b>	The unicast port the node used while in the cluster. This is the port of the node's DatagramSocket for point-to-point communication.
<b>Last Seen</b>	The date and time that the node left the cluster.

## All Services History

Use this display to assess utilization of cache capacity, over time, by all services in a cluster. Analyze load distribution across services and caches, check for bottlenecks and quickly identify services that need more threads. Answer questions such as:

- Is their enough cache capacity available for the service?
- Is their enough storage capacity available for the service?

Use the mouseover tool-tip to see how many caches the service runs on, and data for the selected metric.



**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

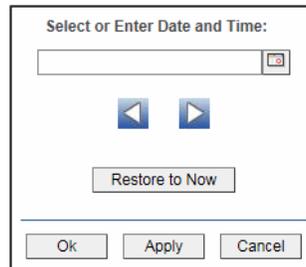
**Filter By:**

The display might include these filtering options:

- Service Metric:** Choose a service metric for which to display data in the heatmap. Use the mouse-over tool-tip to view metrics. Identify a service with high utilization. Perform node analysis by clicking **One** to view the "Single Service History" display.
- CPU%** Percent of CPU utilization in the specified time range.
- Requests** The number of client requests issued to the cluster in the specified time range. This metric is a good indicator of end-user utilization of the service.
- Messages** The number of messages for the given node in the specified time range.
- ActiveThreads** The number of threads in the service thread pool, not currently idle.

- TaskBacklog** The size of the backlog queue that holds tasks scheduled to be executed by one of the service threads. Use this metric for determining capacity utilization for threads running on a service. For example, if the service has a high **TaskBacklog** rate and a low amount of CPU available, consider increasing the number of threads for the service to improve performance.
- RequestPending-Count** The number of pending requests issued by the service.
- RequestAverage-Duration** The average duration (in milliseconds) of an individual request issued by the service since the last time the statistics were reset.

**Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

**Enable MouseOver** Select this option to make service details visible upon mouseover.

**History Heatmap of Selected Metric by Service** Use the heatmap to view utilization trends for all services, over time, and quickly identify heavy usage, indicated by a dark color (by default, dark green). Look for a consistently dark horizontal line, which typically indicates constant high utilization. If this level of utilization is unexpected, consider performing a lower level analysis by viewing service details in the ["Single Service Summary"](#) display.

Two heatmaps, one for Process Nodes and another for Storage Nodes, show utilization trends for the selected metric, for all services running in the cluster. Each row represents a service. Cells in a row are sized uniformly. Each column represents a time period (typically in 10 second intervals). The color of the row cells represent the relative value of the selected service Metric, where a darker shade is a larger value.

Use the mouseover tool-tip to see how many caches the service runs on, and data for the selected metric.

**Services on Process Nodes** Each row represents a service. The color of the cells represents the relative value of the selected Service Metric, where a darker shade is a larger value. The size of the cells are uniform as they each represent one process node. Use the mouseover tool-tip to see how many caches the service runs on, and data for the selected metric.

**Services on Storage Nodes**

Each row represents a service. The color of the cells represents the relative value of the selected Service Metric, where a darker shade is a larger value. The size of the cells are uniform as they each represent one storage node. Use the mouseover tool-tip to see how many caches the service runs on, and data for the selected metric.

**Log Scale**

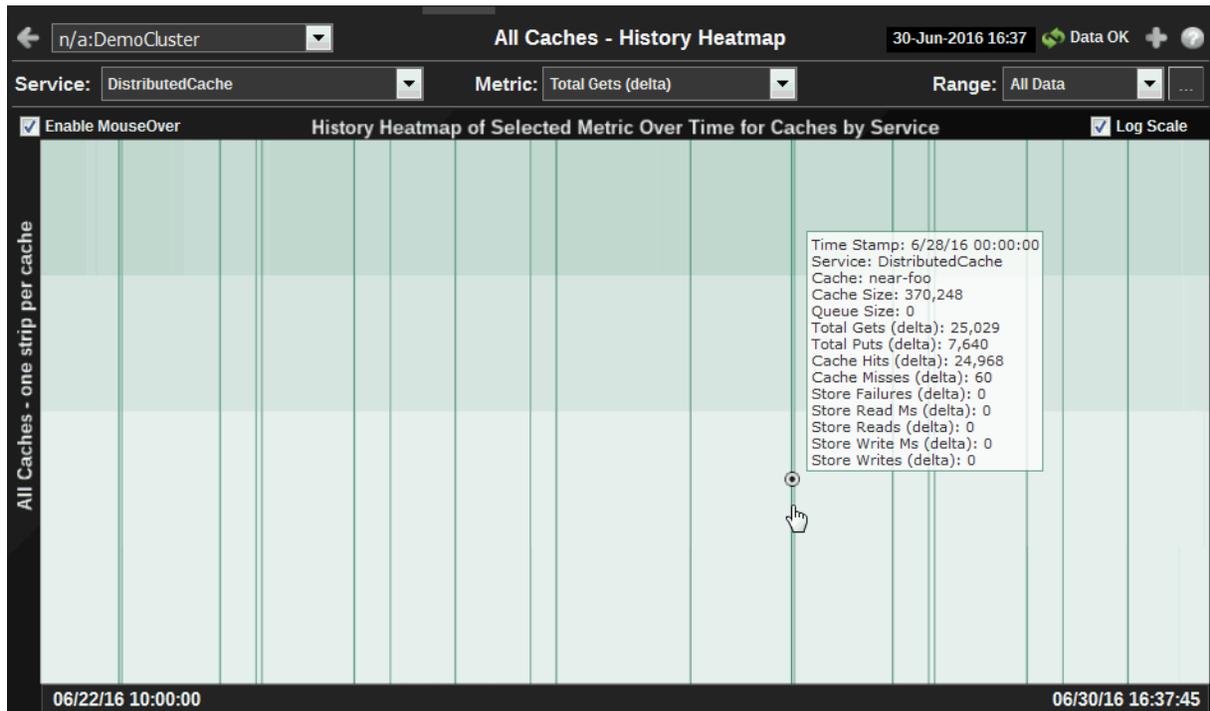
Enable to use a logarithmic scale for the Y axis. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.

**All Caches History**

Use this display to assess capacity utilization, over time, for all caches in a cluster. Analyze load distribution, check for bottlenecks and quickly identify caches with high usage. Answer questions such as:

- Is the cluster using what I expect?
- Is the cluster using it in a uniform scale?

Use the mouseover tool-tip to see the name of the cache and data for the selected metric.



**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** , **Table** open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Filter By:**

- Cluster:** Select a cluster for which to display data in the heatmap.
- Service:** Select a service for which to display data in the heatmap.
- Metric:** Select a metric for which to display data in the heatmap.
- Total Gets** The total number of requests for data from this cache.
- Total Puts** The total number of data stores into this cache.
- Cache Hits** The total number of successful gets for this cache.
- Cache Misses** The total number of failed gets for this cache. This metric indicates whether cache utilization is effective. For example, how often requests are made for data that does not exist in the cache. If a cache has a high rate of misses, consider performing a lower level analysis by viewing the cache in the ["Single Cache Summary"](#) display. Check the metrics for Size, Evictions and Misses to determine whether more capacity is needed.
- Cache Size** The total number of objects in the cache.
- StoreFailures (Delta)** The total number of store failures on this cache since the last data sample.
- StoreReads (Delta)** The total number of load operations on this cache since the last data sample.
- StoreReadMillis (Delta)** The cumulative amount of time (in milliseconds) of load operations for this cache since the last data sample.
- StoreWrites (Delta)** The total number of store and erase operations for this cache since the last data sample.
- StoreWritesMillis (Delta)** The cumulative amount of time (in milliseconds) of store and erase operations on this cache since the last data sample.
- Total Gets** The total number of requests for data from this cache.

- Range:** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .

By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

- AppName:** Choose an AppName to show data for in the display.

**Fields and Data:****AppSlice Information**

<b>Last Update:</b>	The date and time the data was last updated.
<b>Completed:</b>	The total number of completed processes summed across all processes in one AppSlice of the application.
<b>Suspended:</b>	The total number of suspended processes
<b>Failed:</b>	The total number of failed processes
<b>Created Rate:</b>	The number of application processes created per second.
<b>Failed Rate:</b>	The number of failed application processes per second.
<b>Avg Exec:</b>	The average number of seconds for processes to execute.
<b>Avg Elap:</b>	The average amount of elapsed time, in seconds.

**Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

**Enable MouseOver**

Select this option to make cache details visible upon mouseover.

**History Heatmap of Selected Metric**

Use the heatmap to view utilization trends for all caches, over time, and quickly identify heavy usage, indicated by a dark color (by default, dark green). Look for a consistently dark horizontal line, which typically indicates constant high utilization. If this level of utilization is unexpected, consider performing a lower level analysis by viewing cache details in the ["Single Cache Summary"](#) display.

Also look for a dark vertical line, which indicates that all the caches, nodes or services are being used simultaneously. Typically this indicates further analysis is needed.

The heatmap shows cache utilization trends for the selected service and metric, for all caches running in the cluster. Each row represents a cache. Cells in a row are sized uniformly and represent one process node. Each column represents a time period (typically in 10 second intervals). The heatmap is grouped vertically by service. The color of the row cells represent the relative value of the selected service Metric, where a darker shade is a larger value.

Use the mouseover tool-tip to see the name of the cache and data for the selected metric.

**Log Scale** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.

**Base at Zero** Use zero as the Y axis minimum for all graph traces.

**Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar

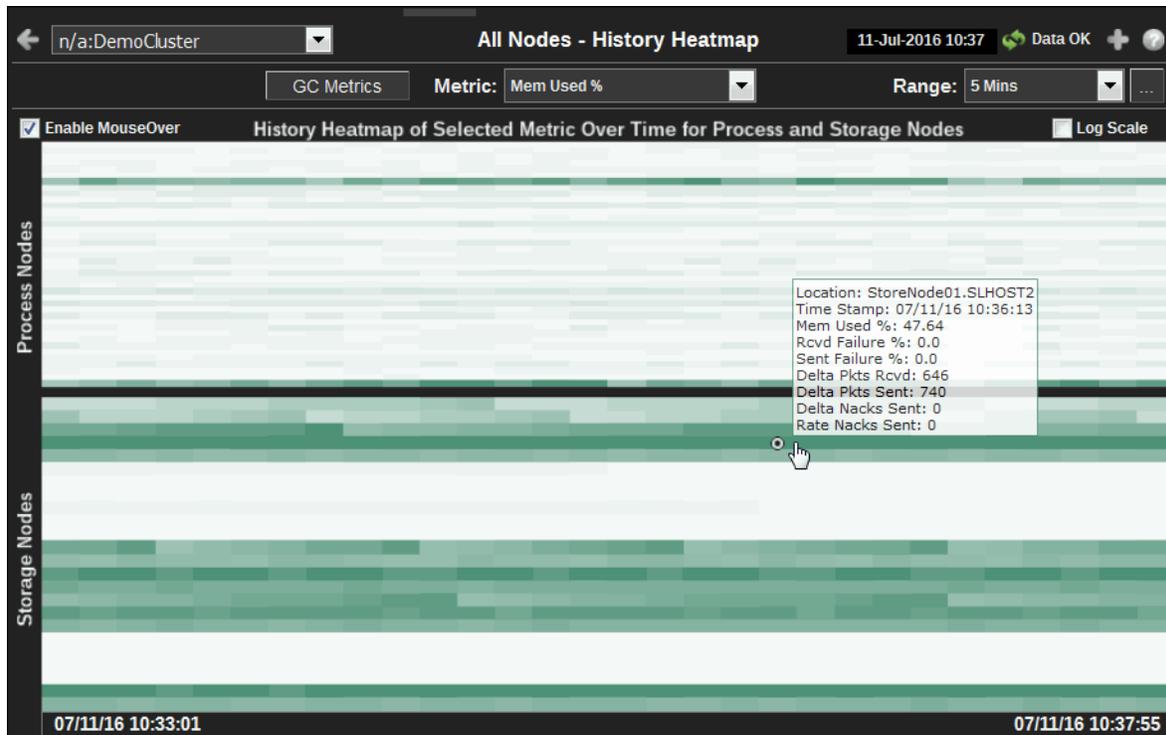
By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## All Nodes History

Use this display to assess capacity utilization, over time, for all nodes in a cluster. Analyze load distribution, check for bottlenecks and quickly identify nodes with high usage. Use the mouseover tool-tip to see the node hostname and data for the selected metric.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

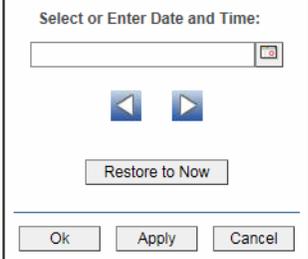
- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

### Filter By:

- Cluster:** Select a cluster for which to display data in the heatmap.
- GC Metrics** Click to open the "All Nodes History" display which shows GC Duty Cycle for all the nodes in a cluster.
- Metric:** Select a metric for which to display data in the heatmap.
  - Mem Used%** The percent (%) of memory used by the node.
  - Packets Sent Fail%** The percent (%) of packets that had to be resent by this node.
  - Packets Rcvd Fail%** The percent (%) of packets that failed to be received by this node.

<b>Delta Packets Sent</b>	The number of packets sent by this node since the last data sample.
<b>Delta Packets Rcvd</b>	The number of packets received by this node since the last data sample.
<b>Delta Nacks Sent</b>	The number of TCMP packets sent by this node since the last data sample. Use this data to troubleshoot communication errors.

**Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

**Enable MouseOver** Select this option to make cache details visible upon mouseover.

**History Heatmap of Selected Metric** Use the heatmap to view utilization trends for all nodes, over time, and quickly identify heavy usage, indicated by a dark color (by default, dark green). Look for a consistently dark horizontal line, which typically indicates constant high utilization. If this level of utilization is unexpected, consider performing a lower level analysis by viewing node details in the ["Node Summary"](#) display.

Two heatmaps, one for Process Nodes and another for Storage Nodes, show utilization trends for the selected metric, for all nodes running in the cluster. Each row represents a node. Cells in a row are sized uniformly. Each column represents a time period (typically in 10 second intervals). The color of the row cells represent the relative value of the selected service Metric, where a darker shade is a larger value.

Use the mouseover tool-tip to see the node hostname and data for the selected metric.

**Process Nodes** Each row represents a node. The color of the cells represents the relative value of the selected Service Metric, where a darker shade is a larger value. The size of the cells are uniform. Use the mouseover tool-tip to see the node hostname and data for the selected metric.

**Storage Nodes**

Each row represents a node. The color of the cells represents the relative value of the selected Service Metric, where a darker shade is a larger value. The size of the cells are uniform. Use the mouseover tool-tip to see the node hostname and data for the selected metric.

**Log Scale**

Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.

## Proxy Services

Proxy Services displays present detailed proxy server performance metrics for the cluster. Use the Proxy Services displays to quickly identify overloaded proxy services and locate the extend client connection causing the issue.

Proxy Services performance metrics include: CPU%, Requests, Request Average Duration, Request Pending Count, Task Backlog and Active Threads.

- [“Proxy / Extend Overview”](#): Heatmap shows the extend connections and a trend graph shows the total connections and total bytes transferred across all proxies for the selected host or hosts.
- [“Proxy / Extend Connections” on page 717](#): Table shows proxy services data with trend graphs/tables of extend connection detail for a specified location.
- [“Proxy / Extend Detail” on page 723](#): Table shows data for proxy services and extend client connection data, including remote endpoint, time stamp, connect time and outgoing byte backlog.
- [“Proxy Nodes History” on page 727](#): Heatmap shows performance utilization, over time, for all proxy service nodes in the selected cluster.
- [“Extend Connections History” on page 729](#): Heatmap shows performance utilization, over time, for all extend connections in the selected cluster.

## Proxy / Extend Overview

Heatmap shows performance utilization and a trend graph shows the total connections and total bytes transferred for all proxy services for the selected host or hosts.



**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** , **Table** open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

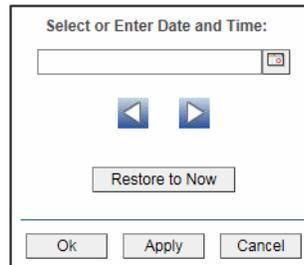
**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

- Cluster:** Select a cluster for which to display data in the heatmap.
- Hosts** Click to open display that shows GC Duty Cycle for all the nodes in a cluster.
- Metric:** Select a metric for which to display data in the heatmap.
  - Alert Severity** The maximum level of alerts in the heatmap rectangle. Values range from 0 - 2, as indicated in the color gradient bar, where 2 is the highest Alert Severity:
  - Alert Count** The total number of critical and warning alerts in the heatmap rectangle. The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.

<b>Current Bytes Sent</b>	Total number of bytes sent by the selected proxy in the time range specified.
<b>Current Bytes Received</b>	Total number of bytes received by the selected proxy in the time range specified.
<b>Proxy CPU%</b>	The average percent CPU utilization for the selected proxy.
<b>Bytes Backlog</b>	The number of pending bytes in the Extend outgoing queue.
<b>Proxy Bytes Backlog</b>	The number of pending bytes in the Proxy outgoing queue.

**Range**

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

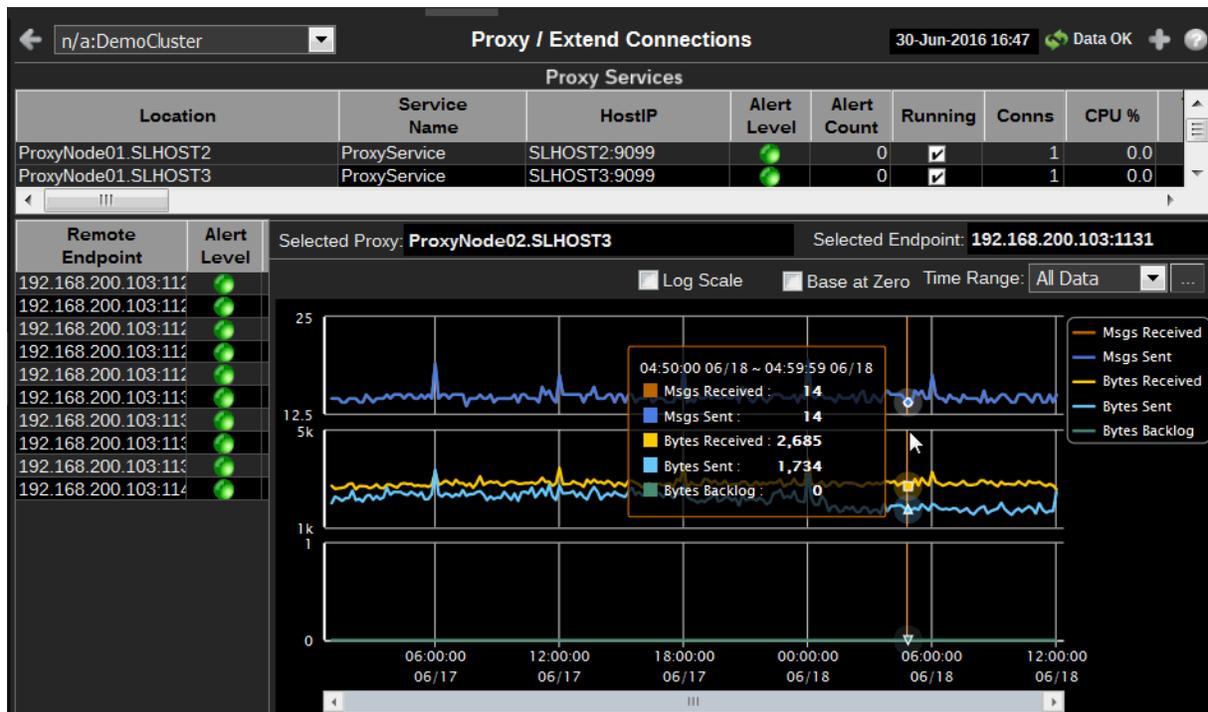
**Total Connections**

The number of extend clients connected to the selected proxy.

**Proxy / Extend Connections**

Table shows proxy services data, including connections, CPU usage and bytes sent and received, and a trend graph displays messages and bytes sent and received for the selected remote endpoint.

The table data is the result of joins of metric from the following Coherence MBeans: Service and ConnectionManager. For details on attributes of these MBeans go to: [http://download.oracle.com/otn\\_hosted\\_doc/coherence/350/com/tangosol/net/management/Registry.html](http://download.oracle.com/otn_hosted_doc/coherence/350/com/tangosol/net/management/Registry.html).



**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu , Table open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

**Filter By:**

The display might include these filtering options:

- Cluster:** Select a cluster from the drop-down menu.
- Proxy Services Location:** A unique identifier for each node. It is defined as: **member\_name.machine.rack.site.**
- HostIP:** The IP address of the host where the proxy service resides.
- Alert Level:** The maximum level of alerts in the row:
  - Red indicates that one or more exceeded their ALARM LEVEL threshold.
  - Yellow indicates that one or more exceeded their WARNING LEVEL threshold.
  - Green indicates that none have exceeded their alert thresholds.

<b>Alert Count</b>	The number of alerts in the row.
<b>Running</b>	Indicates that the proxy service is running when selected.
<b>Connections</b>	The number of extend clients connected to the selected host or hosts.
<b>CPU%</b>	The average percent CPU utilization for each proxy service in the cluster.
<b>Bytes Sent</b>	The number of bytes sent by the proxy service since the proxy service joined the cluster.
<b>Delta</b>	The number of bytes sent by the proxy service since the last data sample.
<b>Backlog</b>	The size (in kilobytes) of the backlog queue.
<b>Bytes Rcvd</b>	The number of bytes received by the proxy service since the proxy service joined the cluster.
<b>Delta</b>	The number of bytes received by the proxy service since the last data sample.
<b>MsgsSent</b>	The number of messages sent by the proxy service since the proxy service joined the cluster.
<b>Delta</b>	The number of messages sent by the proxy service since the last data sample.
<b>Backlog</b>	The size of the backlog queue that holds messages scheduled to be sent by one of the proxy service pool threads.
<b>Msgs Rcvd</b>	The number of messages received by the proxy service since the proxy service joined the cluster.
<b>Delta</b>	The number of messages received by the proxy service since the last data sample.
<b>Tasks</b>	The number of tasks performed by the proxy service since the last time the statistics were reset.
<b>RequestAverageDuration</b>	The average duration (in milliseconds) of an individual synchronous request issued by the proxy service since the last time the statistics were reset.
<b>RequestMaxDuration</b>	Maximum duration (in milliseconds) of an individual proxy service request since the last time the statistics were reset.
<b>RequestTotalCount</b>	The number of requests issued and received by the proxy service.
<b>TaskAverageDuration</b>	The average duration (in milliseconds) of an individual task execution.
<b>TaskBacklog</b>	The size of the backlog queue that holds tasks scheduled to be executed by one of the proxy service pool threads.
<b>TaskCount</b>	The number of tasks performed by the proxy service since the last time the statistics were reset.
<b>TaskHungCount</b>	The total number of currently executing hung tasks.
<b>TaskHungDuration</b>	The longest currently executing hung task duration in milliseconds.
<b>TaskHungTaskId</b>	The id of the of the longest currently executing hung task.

<b>TaskHungThresholdMillis</b>	The duration (in milliseconds) that a proxy service task can execute before it is considered hung. Note that a posted task that has not yet started is never considered as hung.
<b>TaskMaxBacklog</b>	The maximum size of the proxy service backlog queue since the last time the statistics were reset.
<b>TaskTimeoutCount</b>	The total number of timed-out proxy service tasks since the last time the statistics were reset.
<b>RequestPendingCount</b>	The number of pending proxy service requests.
<b>RequestPendingDuration</b>	The average duration (in milliseconds) that an individual proxy service request waits before being executed.
<b>RequestTimeoutCount</b>	The total number of timed-out proxy service requests since the last time the statistics were reset.
<b>RequestTimeoutMillis</b>	The duration (in milliseconds) for a proxy service request to reach the specified timeout threshold.
<b>TaskTimeoutMillis</b>	The default timeout value (in milliseconds) for tasks that can be timed-out but do not explicitly specify the task execution timeout value.
<b>IncomingBufferPoolSize</b>	<b>The number of buffers in the incoming pool.</b>
<b>ThreadAbandonedCount</b>	The number of abandoned threads from the proxy service thread pool. A thread is abandoned and replaced with a new thread if it executes a task for a period of time longer than execution timeout and all attempts to interrupt it fail.
<b>ThreadCount</b>	The number of threads in the proxy service thread pool.
<b>ThreadIdleCount</b>	The number of currently idle threads in the proxy service thread pool.
<b>AverageActiveThreadCount</b>	The average number of proxy service active threads, not currently idle, since the last time the statistics were reset.
<b>ThreadAverageActiveCount</b>	<b>The average number of active (not idle) threads in the service thread pool since the last time the statistics were reset.</b>
<b>AverageTaskDuration</b>	The average duration (in milliseconds) to perform a proxy service task since the last time the statistics were reset.
<b>MaximumBacklog</b>	The maximum size of the backlog queue since the last time the statistics were reset.
<b>Throughput</b>	The amount of data (in kilobytes) that is transferred by the proxy service.
<b>ThroughputInbound</b>	The amount of data (in kilobytes) that is transferred from clients to the proxy service.
<b>ThroughputOutbound</b>	The amount of data (in kilobytes) that is transferred from the proxy service to clients.
<b>IncomingBufferPoolCapacity</b>	The size (in kilobytes) of the proxy service incoming buffer pool.
<b>OutgoingBufferPoolCapacity</b>	The size (in kilobytes) of the proxy service outgoing buffer pool.
<b>OutgoingBufferPoolSize</b>	The number of buffers in the proxy service outgoing pool.
<b>nodeld</b>	The unique identifier for the proxy service.

	<b>RefreshTime</b>	The timestamp when this model was last retrieved from a corresponding node. For local servers it is the local time.
	<b>HostName</b>	The name of the host where the proxy service resides.
	<b>MemberName</b>	A specified, unique name of the host where the proxy service resides.
	<b>SeniorMemberId</b>	The proxy service senior member id. If the proxy service is not running, it is -1.
	<b>Rate</b>	The number of errors accumulated per second.
<b>Execution</b>	<b>Min</b>	The shortest execution time of any process instance, in milliseconds.
	<b>Max</b>	The longest execution time of any process instance, in milliseconds.
	<b>Average</b>	The average execution time for all completed process instances, in milliseconds.
	<b>Current</b>	The amount of time accumulated this update cycle.
	<b>Rate</b>	The amount of time accumulated per second.
<b>Elapsed</b>	<b>Min</b>	The shortest elapsed time of any process instance, in milliseconds.
	<b>Max</b>	The longest elapsed time of any process instance, in milliseconds.
	<b>Average</b>	The average elapsed time for all completed process instances, in milliseconds.
	<b>Current</b>	The amount of elapsed time accumulated this update cycle.
	<b>Rate</b>	The amount of elapsed time accumulated per second.
<b>Selected Proxy</b>	<b>This field is populated by the selection made in the Proxy Services table.</b>	
<b>Selected Endpoint</b>	This field is populated by the selection made in the Remote Endpoint table.	

### Trend Graphs

Select a host from the Proxy Services table and a connection from the Remote Endpoint table. This table is populated by the selection made in the Proxy Services table.

Alert Level shows the maximum level of alerts in row:

- Red indicates that one or more exceeded their ALARM LEVEL threshold.
- Yellow indicates that one or more exceeded their WARNING LEVEL threshold.
- Green indicates that none have exceeded their alert thresholds.

**Msgs Received:** Traces the number of messages received by the selected proxy service from the remote endpoint.

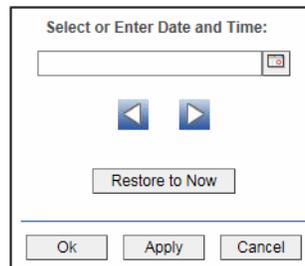
**Msgs Sent:** Traces the number of bytes received by the selected proxy service from the remote endpoint.

**Bytes Received:** Traces the rate at which the application is accumulating process execution time, in milliseconds per second.

**Bytes Sent:** Traces the number of executed activities per second.

**All Activities Exec Time/sec:** Traces the number of bytes sent by the selected proxy service to the remote endpoint.

**Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Proxy / Extend Detail

Table shows data for proxy services and extend client connection data, including remote endpoint, time stamp, connect time and outgoing byte backlog.

Proxy Services								
Location	Service Name	HostIP	Alert Level	Alert Count	Running	Conns	CPU %	Tot
ProxyNode01.SLHOST2	ProxyService	SLHOST2:9099	🟢	0	☑	1	0.0	2
ProxyNode01.SLHOST3	ProxyService	SLHOST3:9099	🟢	0	☑	1	0.0	2
ProxyNode01.SLHOST4	ProxyService	SLHOST4:9099	🟢	0	☑	0	0.0	4
ProxyNode02.SLHOST2	ProxyService	SLHOST2:9099	🟢	0	☑	0	0.0	
ProxyNode02.SLHOST3	ProxyService	SLHOST3:9099	🟢	0	☑	0	0.0	
ProxyNode02.SLHOST4	ProxyService	SLHOST4:9099	🟢	0	☑	0	0.0	

Extend Client Connections for ProxyNode02.SLHOST3:*							
Location	Remote Endpoint	Alert Level	Alert Count	Time Stamp	Connect Time (mins)	Bytes Backlog	Message Backlog
ProxyNode02.SLHOST3	192.168.200.103:1131	🟢	0	06/17/16 01:03:29	2,098.6	0	
ProxyNode02.SLHOST3	192.168.200.103:1136	🟢	0	06/17/16 01:03:34	2,098.5	0	
ProxyNode02.SLHOST3	192.168.200.103:1128	🟢	0	06/19/16 01:03:33	659.1	0	
ProxyNode02.SLHOST3	192.168.200.103:1121	🟢	0	06/19/16 01:03:27	1,437.2	0	
ProxyNode02.SLHOST3	192.168.200.103:1124	🟢	0	06/22/16 01:03:32	2,818.6	0	
ProxyNode02.SLHOST3	192.168.200.103:1132	🟢	0	06/22/16 01:03:37	2,098.7	0	
ProxyNode02.SLHOST3	192.168.200.103:1136	🟢	0	06/24/16 01:03:40	1,018.4	0	
ProxyNode02.SLHOST3	192.168.200.103:1140	🟢	0	06/24/16 01:03:45	1,018.3	0	
ProxyNode02.SLHOST3	192.168.200.103:1121	🟢	0	06/26/16 01:03:29	1,437.2	0	
ProxyNode02.SLHOST3	192.168.200.103:1127	🟢	0	06/26/16 01:03:34	1,437.1	0	

### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

**Cluster** Select a cluster from the drop-down menu.

### Proxy Services

<b>Location</b>	A unique identifier for each node. It is defined as: <b>member_name.machine.rack.site</b> .
<b>HostIP</b>	The IP address of the host where the proxy service resides.
<b>Running</b>	Indicates that the proxy service is running when selected.
<b>Connections</b>	The number of extend clients connected to the selected host or hosts.
<b>CPU%</b>	The average percent CPU utilization for each proxy service in the cluster.

<b>Bytes Sent</b>	The number of bytes sent by the proxy service since the proxy service joined the cluster.
<b>Delta</b>	The number of bytes sent by the proxy service since the last data sample.
<b>Backlog</b>	The size (in kilobytes) of the backlog queue.
<b>Bytes Rcvd</b>	The number of bytes received by the proxy service since the proxy service joined the cluster.
<b>Delta</b>	The number of bytes received by the proxy service since the last data sample.
<b>MsgsSent</b>	The number of messages sent by the proxy service since the proxy service joined the cluster.
<b>Delta</b>	The number of messages sent by the proxy service since the last data sample.
<b>Backlog</b>	The size of the backlog queue that holds messages scheduled to be sent by one of the proxy service pool threads.
<b>Msgs Rcvd</b>	The number of messages received by the proxy service since the proxy service joined the cluster.
<b>Delta</b>	The number of messages received by the proxy service since the last data sample.
<b>Tasks</b>	The number of tasks performed by the proxy service since the last time the statistics were reset.
<b>RequestAverageDuration</b>	The average duration (in milliseconds) of an individual synchronous request issued by the proxy service since the last time the statistics were reset.
<b>RequestMaxDuration</b>	Maximum duration (in milliseconds) of an individual proxy service request since the last time the statistics were reset.
<b>RequestPendingCount</b>	The number of pending proxy service requests.
<b>RequestPendingDuration</b>	The average duration (in milliseconds) that an individual proxy service request waits before being executed.
<b>RequestTimeoutCount</b>	The total number of timed-out proxy service requests since the last time the statistics were reset.
<b>RequestTimeoutMillis</b>	The duration (in milliseconds) for a proxy service request to reach the specified timeout threshold.
<b>RequestTotalCount</b>	The number of requests issued and received by the proxy service.
<b>TaskAverageDuration</b>	The average duration (in milliseconds) of an individual task execution.
<b>TaskBacklog</b>	The size of the backlog queue that holds tasks scheduled to be executed by one of the proxy service pool threads.
<b>TaskCount</b>	The number of tasks performed by the proxy service since the last time the statistics were reset.
<b>TaskHungCount</b>	The total number of currently executing hung tasks.
<b>TaskHungDuration</b>	The longest currently executing hung task duration in milliseconds.
<b>TaskHungTaskId</b>	The id of the of the longest currently executing hung task.

<b>TaskHungThresholdMillis</b>	The duration (in milliseconds) that a proxy service task can execute before it is considered hung. Note that a posted task that has not yet started is never considered as hung.
<b>TaskMaxBacklog</b>	The maximum size of the proxy service backlog queue since the last time the statistics were reset.
<b>TaskTimeoutCount</b>	The total number of timed-out proxy service tasks since the last time the statistics were reset.
<b>TaskTimeoutMillis</b>	The default timeout value (in milliseconds) for tasks that can be timed-out but do not explicitly specify the task execution timeout value.
<b>IncomingBufferPoolSize</b>	The number of buffers in the incoming pool.
<b>ThreadAbandonedCount</b>	The number of abandoned threads from the proxy service thread pool. A thread is abandoned and replaced with a new thread if it executes a task for a period of time longer than execution timeout and all attempts to interrupt it fail.
<b>ThreadCount</b>	The number of threads in the proxy service thread pool.
<b>ThreadIdleCount</b>	The number of currently idle threads in the proxy service thread pool.
<b>AverageActiveThreadCount</b>	The average number of proxy service active threads, not currently idle, since the last time the statistics were reset.
<b>ThreadAverageActiveCount</b>	The average number of active (not idle) threads in the service thread pool since the last time the statistics were reset.
<b>AverageTaskDuration</b>	The average duration (in milliseconds) to perform a proxy service task since the last time the statistics were reset.
<b>MaximumBacklog</b>	The maximum size of the backlog queue since the last time the statistics were reset.
<b>Throughput</b>	The amount of data (in kilobytes) that is transferred by the proxy service.
<b>ThroughputInbound</b>	The amount of data (in kilobytes) that is transferred from clients to the proxy service.
<b>ThroughputOutbound</b>	The amount of data (in kilobytes) that is transferred from the proxy service to clients.
<b>IncomingBufferPoolCapacity</b>	The size (in kilobytes) of the proxy service incoming buffer pool.
<b>OutgoingBufferPoolCapacity</b>	The size (in kilobytes) of the proxy service outgoing buffer pool.
<b>OutgoingBufferPoolSize</b>	The number of buffers in the proxy service outgoing pool.
<b>nodeId</b>	The unique identifier for the proxy service.
<b>RefreshTime</b>	The timestamp when this model was last retrieved from a corresponding node. For local servers it is the local time.
<b>HostName</b>	The name of the host where the proxy service resides.
<b>MemberName</b>	A specified, unique name of the host where the proxy service resides.
<b>SeniorMemberId</b>	The proxy service senior member id. If the proxy service is not running, it is -1.

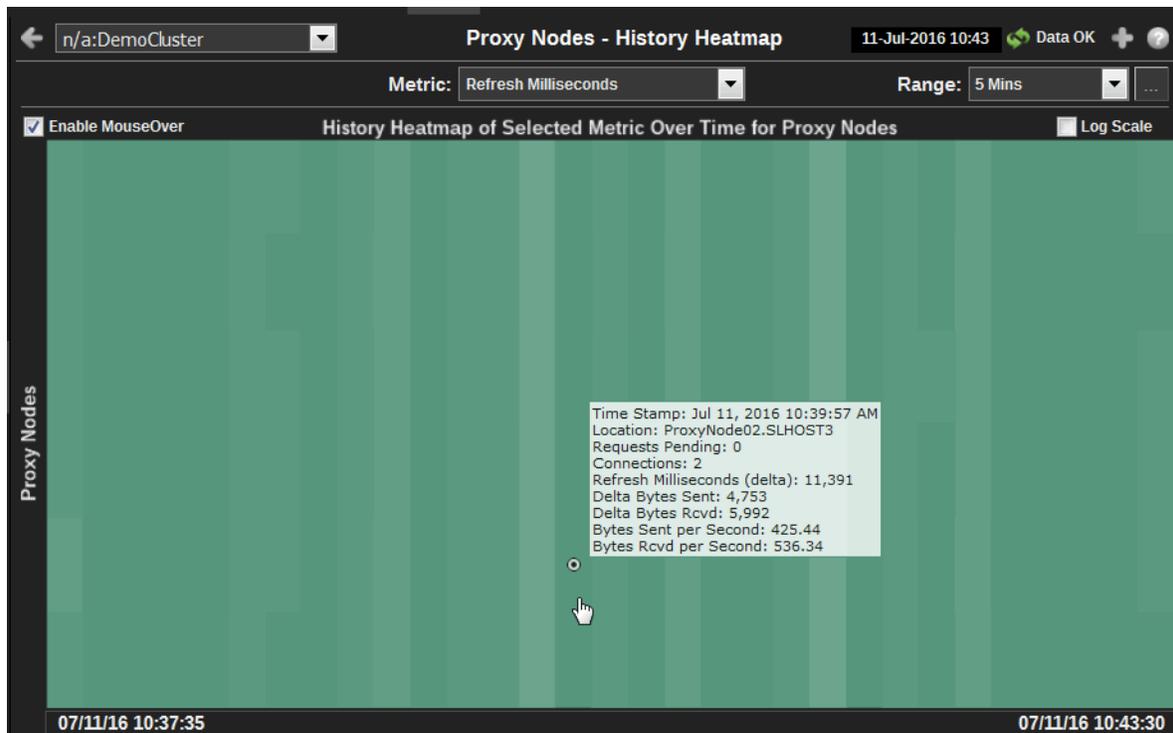
**Extend  
Client  
Connections**

Select a row from the Proxy Services table to populate client data in the table.

<b>Location</b>	A unique identifier for each node. It is defined as: <b>member_name.machine.rack.site.</b>
<b>RemoteEndpoint</b>	The IP address of the client.
<b>Timestamp</b>	The date and time (in cluster time) that this client joined the proxy service.
<b>Connect Time (mins)</b>	The duration (in minutes) the client has been connected to the proxy service.
<b>OutgoingByteBacklog</b>	The size of the backlog queue (in bytes) that holds outgoing bytes scheduled to be executed by one of the proxy service pool threads for the client.
<b>OutgoingMessageBacklog</b>	The number of messages in the backlog queue that holds outgoing messages scheduled to be sent to the client by one of the proxy service pool threads.
<b>TotalBytesReceived</b>	The number of bytes received from the client by the proxy service since the client connected to the proxy service.
<b>Delta</b>	The number of bytes received from the client by the proxy service since the last data sample.
<b>TotalBytesSent</b>	The number of bytes sent to the client by the proxy service since the client connected to the proxy service.
<b>Delta</b>	The number of bytes sent to the client by the proxy service since the last data sample.
<b>TotalMessagesReceived</b>	The number of messages received from the client by the proxy service since the client connected to the proxy service.
<b>Delta</b>	The number of messages received from the client by the proxy service since the last data sample.
<b>TotalMessagesSent</b>	The number of messages sent to the client by the proxy service since the client connected to the proxy service.
<b>Delta</b>	The number of messages sent to the client by the proxy service since the last data sample.
<b>UUID</b>	The unique identifier for the extend client application.
<b>nodeld</b>	The unique identifier for the proxy service the extend client is connected to.

## Proxy Nodes History

Heatmap shows performance utilization, over time, for all proxy service nodes in the selected cluster. Use this display to assess performance, over time, for all proxy service nodes in a cluster. Analyze load distribution, check for bottlenecks and quickly identify proxy service nodes with high usage.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Cluster** Select a cluster from the drop-down menu.

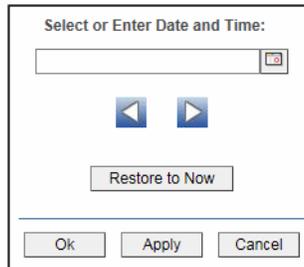
**Metric** Select a metric from the drop-down menu.

- Request Pending** The number of pending requests issued by the node.
- Connections** Total number of connection for the node.
- Refresh Milliseconds** The amount of time, in milliseconds, since the last data sample.
- Delta Bytes Sent** Total number of bytes sent by the node since the last data sample.
- Delta Bytes Rcvd** Total number of bytes received by the node since the last data sample.

**Bytes Sent Per Second** Total bytes sent, per second, by the node.

**Range**

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

**Enable MouseOver**

Select this option to make details visible upon mouseover.

**Log Scale**

Enable to use a logarithmic scale for the Y axis. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.

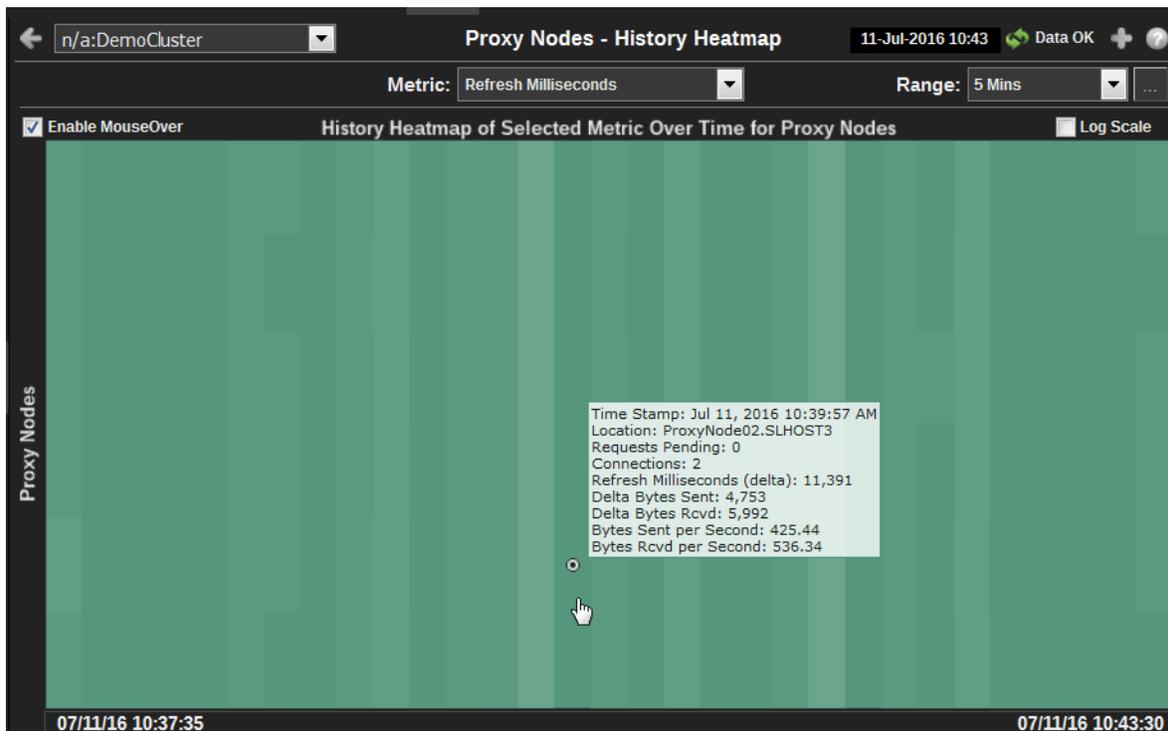
**Proxy Nodes Heatmap**

Use the heatmap to view utilization trends for all Process and Storage nodes, over time, and quickly identify heavy usage, indicated by a dark color (by default, dark green). Look for a consistently dark horizontal line, which typically indicates constant high utilization. If this level of utilization is unexpected, consider performing a lower level analysis by viewing node details in the ["Node Summary"](#) display.

Use the mouseover tool-tip to see the node hostname and data for the selected metric.

## Extend Connections History

Heatmap shows performance utilization, over time, for all extend connections in the selected cluster.



**Title Bar** (possible features are):

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu, Table open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

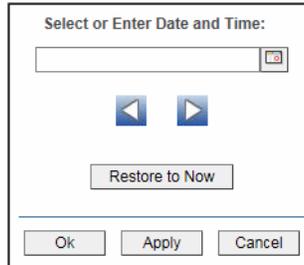
**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

<b>Cluster</b>	Select a cluster from the drop-down menu.
<b>Metric</b>	Select a metric from the drop-down menu.
<b>Delta Bytes Sent</b>	Total number of bytes sent by the node since the last data sample.
<b>Delta Bytes Rcvd</b>	Total number of bytes received by the node since the last data sample.
<b>Delta Messages Sent</b>	Total number of messages sent by the node since the last data sample.
<b>Delta Messages Rcvd</b>	Total number of messages received by the node since the last data sample.
<b>Bytes Sent per Second</b>	Total bytes sent, per second, by the node.

- Bytes per Second** Total bytes received, per second, by the node.
- Msgs Sent per Second** Total messages sent, per second, by the node.
- Msgs Rcvd per Second** Total messages received, per second, by the node.

**Range**

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

**Enable MouseOver**

Select this option to make details visible upon mouseover.

**Log Scale**

Enable to use a logarithmic scale for the Y axis. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.

**Proxy Nodes Heatmap**

Use the heatmap to view utilization trends for all Process and Storage nodes, over time, and quickly identify heavy usage, indicated by a dark color (by default, dark green). Look for a consistently dark horizontal line, which typically indicates constant high utilization. If this level of utilization is unexpected, consider performing a lower level analysis by viewing node details in the Single Node - Summary display.

Use the mouseover tool-tip to see the node hostname and data for the selected metric.

## Cache Services

Cache Services displays present detailed service performance metrics for the cluster. Use the Cache Services displays to quickly identify overloaded services and locate the client connection causing the issue.

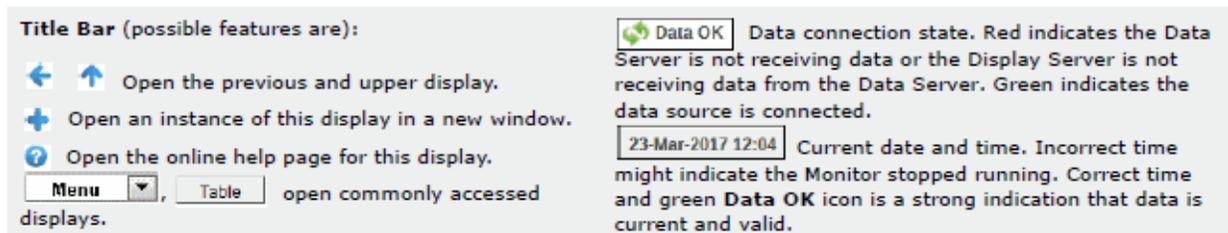
These displays show metrics for all cache services, including: CPU%, Requests, Request Average Duration, Request Pending Count, Task Backlog and Active Threads.

- ["Single Service Summary" on page 731](#): Trend graphs show performance metrics for a single service aggregated across all nodes.
- ["Service Metrics Overview" on page 735](#): Heatmap shows overview of the current behavior of the cluster, displaying metrics across nodes in the cluster for a selected service or for all services. Enables you to determine if the behavior of the cluster is balanced across all nodes or identify if some nodes are hot spots.
- ["Service Metric Heatmap" on page 736](#): Heatmap shows current value of a selected metric, selected by service, across the cluster. Enables you to determine if the behavior of the cluster, for the selected metric, is balanced or identify if some nodes are hot spots.
- ["Single Service History" on page 738](#): Use this display to perform low-level analysis of service capacity utilization, over time, per node. Heatmap enables you to view the impact of events across the cluster as well as the relative historical performance of nodes across the cluster.
- ["Cache Service Detail" on page 740](#): Table view of attributes of a selected service for a selected host for nodes. Attribute values can be ordered to identify the nodes with the highest and lowest values of interest.

## Single Service Summary

This display shows performance metrics for a single service aggregated across all nodes.





<b>Cluster</b>	Select a cluster to display.
<b>Service</b>	Select a service to display.
<b>Storage Nodes</b>	Select to display storage node data in the trend graphs of this display.
<b>Process Nodes</b>	Select to display process node data in the trend graphs of this display.
<b>Caches</b>	The number of caches managed by the service.
<b>Type</b>	The type of cache.
<b>Storage Nodes</b>	The number of storage nodes in the cache.
<b>Process Nodes</b>	The number of process nodes in the cache.
<b>Status</b>	<p>The high availability status of the service:</p> <ul style="list-style-type: none"> <li>● <b>ENDANGERED:</b> There is potential data loss in the cluster if a node goes offline.</li> <li>● <b>NODE-SAFE:</b> There is no risk of data loss in the cluster if a node goes offline (or is taken offline using kill-9). The data is replicated across multiple nodes and remains available in the cluster.</li> <li>● <b>MACHINE-SAFE:</b> There is no risk of data loss in the cluster if a machine goes offline (or is taken offline using kill-9). The data is replicated across multiple machines and remains available in the cluster.</li> <li>● <b>RACK-SAFE:</b> There is no risk of data loss in the cluster if a rack goes offline (or is taken offline using kill-9). The data is replicated across multiple racks and remains available in the cluster.</li> <li>● <b>SITE-SAFE:</b> There is no risk of data loss in the cluster if a site goes offline (or is taken offline using kill-9). The data is replicated across multiple sites and remains available in the cluster.</li> </ul>
<b>Requests</b>	<p>Requests executed by the service.</p> <p><b>Total</b> The number of requests executed.</p> <p><b>Rate / Delta</b> Use the <b>Use Rates</b> checkbox to toggle between two value types: <b>Rate</b> and Delta (as labeled in the display upon selection).</p> <p>When the <b>Use Rates</b> (checkbox) is NOT selected the Delta values are shown here and in the trend graphs. Delta is the difference in the value since the last sample. When the <b>Use Rates</b> (checkbox) is selected the Rate values are shown here and in the trend graphs. Rate is the value per second. The Rate value is useful when the sampling time period is unknown, has changed, or has a long duration specified. For a given rate, the Rate value does not vary if the sample period changes (whereas the Delta value does vary). The Rate value enables you to directly compare rates on systems with different sample periods.</p> <p><b>Pending</b> The number of pending requests.</p>

**Messages**

Messages executed by the service.

**Total** The number of messages executed.

**Rate / Delta** Use the **Use Rates** checkbox to toggle between two value types: **Rate** and **Delta** (as labeled in the display upon selection).

When the **Use Rates** (checkbox) is NOT selected the **Delta** values are shown here and in the trend graphs. **Delta** is the difference in the value since the last sample. When the **Use Rates** (checkbox) is selected the **Rate** values are shown here and in the trend graphs. **Rate** is the value per second. The **Rate** value is useful when the sampling time period is unknown, has changed, or has a long duration specified. For a given rate, the **Rate** value does not vary if the sample period changes (whereas the **Delta** value does vary). The **Rate** value enables you to directly compare rates on systems with different sample periods.

**Req Avg Duration** The average amount of time to process messages.

**Tasks**

Tasks performed by the service.

**Count** The number of tasks performed.

**Backlog** The number of tasks scheduled to be executed by one of the service threads.

**Queue** The Write Back Queue total across all caches on the service.

**Threads**

Threads on the service.

**Count** The number of threads on the service.

**Active** The number of threads in the service not currently idle.

**Avg CPU%** The average amount of CPU usage (%) for the service.

**Storage /  
Process Node  
Totals**

The trend graphs show aggregated performance metrics for storage or process nodes. Choose **Storage Nodes** or **Process Nodes** at the top of this display.

**Use Rates** Select to show **Rate** values in the **Requests and Messages** fields and trend graphs.

**Rate** is the value per second. The **Rate** value is useful when the sampling time period is unknown, has changed, or has a long duration specified. For a given rate, the **Rate** value does not vary if the sample period changes (whereas the **Delta** value does vary). The **Rate** value enables you to directly compare rates on systems with different sample periods. Deselect **Use Rates** to show the **Delta** values in the **Activity - Current (Delta)** fields and trend graphs. **Delta** is the difference in the value since the last sample.

**Log Scale** Enable to use a logarithmic scale for the Y axis. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.

**Base at Zero** Use zero for the Y axis minimum for all graphs.

**Time Range**

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .

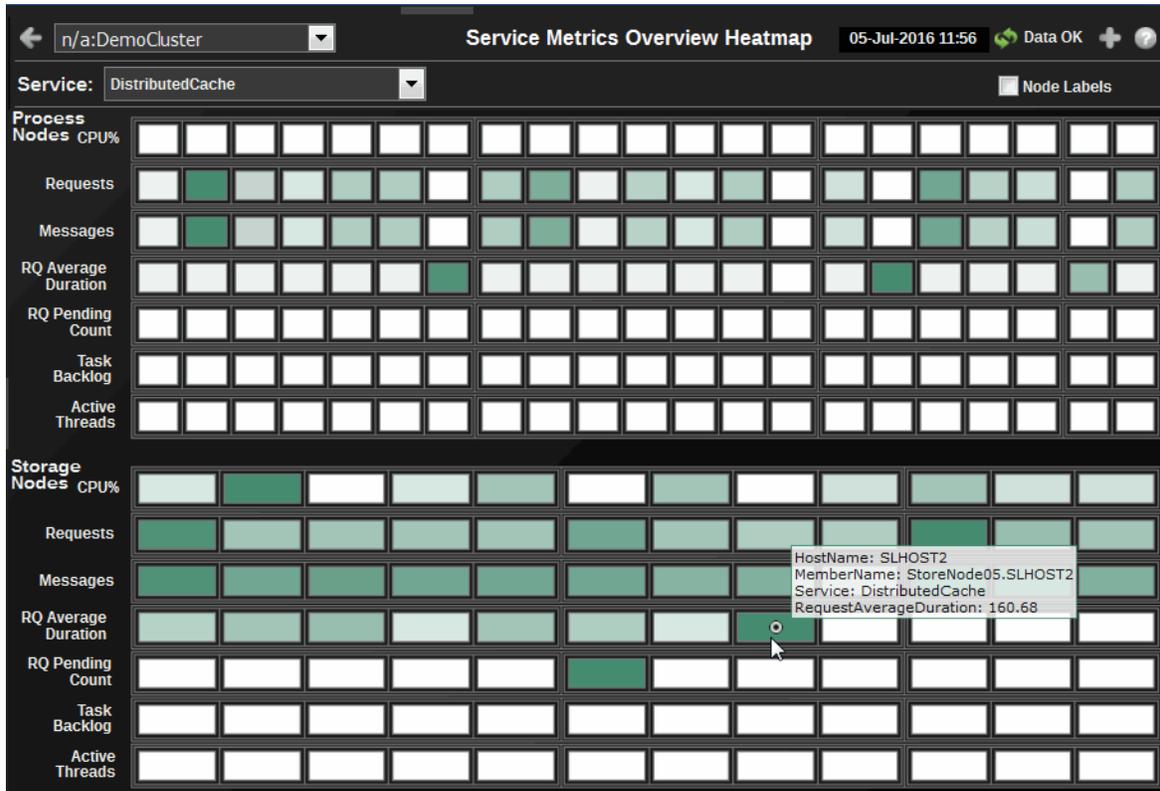
By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Service Metrics Overview

Heatmap of Process (non-storage enabled) Nodes and Storage (enabled) Nodes. Size = One Node. Color = Relative Value of Selected Metric.



**Title Bar** (possible features are):

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu, Table open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

**Cluster** Select a cluster to display.

**Service** Select a service to display.

**Node Labels** Select to display node labels.

**Process Nodes**  
**Storage Nodes** Color of the cells represents the relative value of the selected Metric; a darker shade is a larger value. The size of all cells is identical as they each represent one process node.

**CPU%** Percent of CPU utilization on the given node.

**Requests** Number of requests issued by the service in the measured period.

**Messages** The number of messages for the given node in the measured interval.

**Request Average Duration** Average duration (in milliseconds) of an individual request issued by the service since the last time the statistics were reset.

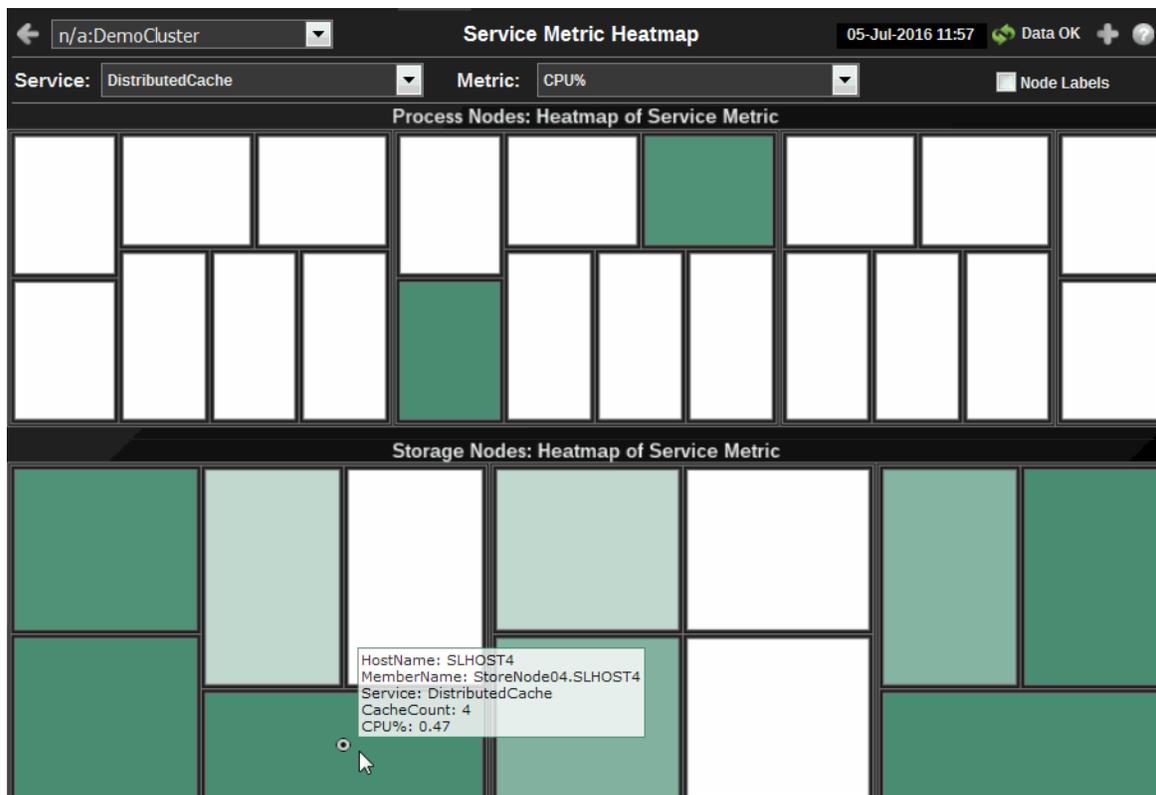
**Request Pending Count** Number of pending requests issued by the service.

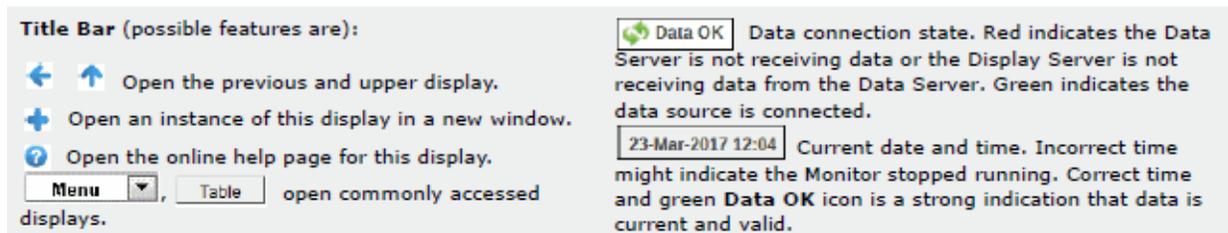
**Task Backlog** Size of the backlog queue that holds tasks scheduled to be executed by one of the service threads.

**Active Threads** Number of threads in the service thread pool, not currently idle.

### Service Metric Heatmap

Heatmap of Process (non-storage enabled) Nodes and Storage (enabled) Nodes. Size = Number of Caches in Selected Service, Color = Relative Value of Selected Metric.





<b>Cluster</b>	Select a cluster to display.
<b>Service</b>	Select a service to display.
<b>Node Labels</b>	Select to display node labels.
<b>Metric</b>	<p><b>CPU%</b> Percent of CPU utilization on the given node.</p> <p><b>Requests</b> Number of requests issued by the service in the measured period.</p> <p><b>Request Average Duration</b> Average duration (in milliseconds) of an individual request issued by the service since the last time the statistics were reset.</p> <p><b>Request Pending Count</b> Number of pending requests issued by the service.</p> <p><b>Task Backlog</b> Size of the backlog queue that holds tasks scheduled to be executed by one of the service threads.</p> <p><b>Active Threads</b> Number of threads in the service thread pool, not currently idle.</p>
<b>Node Labels</b>	Select to view node locations. <b>Location</b> is a unique identifier for each node and defined as: <b>member_name.machine.rack.site</b> .
<b>Process Nodes: Heatmap of Service Metric</b>	<p>Color of the cells represents the relative value of the selected Metric for a given process node; a darker shade is a larger value.</p> <p>Size of the cells is based the number of caches in the selected Service for that process node.</p>
<b>Storage Nodes: Heatmap of Service Metric</b>	<p>Color of the cells represents the relative value of the selected Metric for a given process node; a darker shade is a larger value.</p> <p>Size of the cells is based the number of caches in the selected Service for that process node.</p>

## Single Service History

Use this display to perform low-level analysis, node-by-node, of service capacity utilization. Heatmap of Process (non storage enabled) Nodes and Storage (enabled) Nodes. Color = Relative Value of Selected Metric.



### Title Bar (possible features are):

-   Open the previous and upper display.
-  Open an instance of this display in a new window.
-  Open the online help page for this display.
- open commonly accessed displays.

 **Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

- Cluster** Select a cluster to display.
- All** Click to view the ["All Services History"](#) display.
- Service** Select a service to display.

**Metric**

**CPU%** CPU Utilization (as a percent) on the given node.

**Requests** Number of requests issued by the service in the measured period.

**Request Average Duration** Average duration (in milliseconds) of an individual request issued by the service since the last time the statistics were reset.

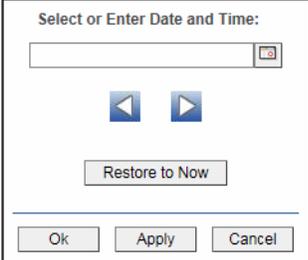
**Request Pending Count** Number of pending requests issued by the service.

**Task Backlog** Size of the backlog queue that holds tasks scheduled to be executed by one of the service threads.

**Active Threads** Number of threads in the service thread pool, not currently idle.

**Time Range**

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

**Process Nodes:  
History Heatmap  
of Service Metric**

Color of the cells represents the relative value of the selected Metric for a given process node; a darker shade is a larger value.

The value of the Metric is displayed over the specified History for all process nodes in the selected Service.

**Storage Nodes:  
History Heatmap  
of Service Metric**

Color of the cells represents the relative value of the selected Metric for a given storage node; a darker shade is a larger value.

The value of the Metric is displayed over the specified History for all storage nodes in the selected Service.

## Cache Service Detail

This display provides a table view of attributes of a selected service for a selected host for nodes. Attribute values can be ordered to identify the nodes with the highest and lowest values of interest.

Location	Service	Running	StatusHA	Storage	CPU %	Messages
JmxNode01.SLHOST1	DistributedCache	☑	MACHINE-SAFE	☐	0.0	195,635
ProcessNode01.SLHOST2	DistributedCache	☑	MACHINE-SAFE	☐	0.0	376,689
ProcessNode01.SLHOST3	DistributedCache	☑	MACHINE-SAFE	☐	0.0	1,149,263
ProcessNode01.SLHOST4	DistributedCache	☑	MACHINE-SAFE	☐	0.0	1,131,618
ProcessNode04.SLHOST2	DistributedCache	☑	MACHINE-SAFE	☐	0.0	188,285
ProcessNode04.SLHOST3	DistributedCache	☑	MACHINE-SAFE	☐	0.0	574,088
ProcessNode04.SLHOST4	DistributedCache	☑	MACHINE-SAFE	☐	0.0	564,263
ProcessNode05.SLHOST3	DistributedCache	☑	MACHINE-SAFE	☐	0.0	1,150,706
ProcessNode05.SLHOST4	DistributedCache	☑	MACHINE-SAFE	☐	0.0	1,134,618
ProcessNode05n.SLHOST2	DistributedCache	☑	MACHINE-SAFE	☐	0.0	754,247
ProcessNode05n.SLHOST3	DistributedCache	☑	MACHINE-SAFE	☐	0.0	2,334,486
ProcessNode05n.SLHOST4	DistributedCache	☑	MACHINE-SAFE	☐	0.0	2,248,394
ProcessNode08.SLHOST3	DistributedCache	☑	MACHINE-SAFE	☐	0.0	573,289
ProcessNode08.SLHOST4	DistributedCache	☑	MACHINE-SAFE	☐	0.0	565,373
ProxyNode01.SLHOST2	DistributedCache	☑	MACHINE-SAFE	☐	0.0	375,960
ProxyNode01.SLHOST3	DistributedCache	☑	MACHINE-SAFE	☐	0.0	2,431
ProxyNode01.SLHOST4	DistributedCache	☑	MACHINE-SAFE	☐	0.0	1,788
ProxyNode02.SLHOST2	DistributedCache	☑	MACHINE-SAFE	☐	0.0	307
ProxyNode02.SLHOST3	DistributedCache	☑	MACHINE-SAFE	☐	0.0	584,044
ProxyNode02.SLHOST4	DistributedCache	☑	MACHINE-SAFE	☐	0.0	1,127,855
RTViewOCDataServer.SLHOST1	DistributedCache	☑	MACHINE-SAFE	☐	0.0	92
StoreNode01.SLHOST1	DistributedCache	☑	MACHINE-SAFE	☑	0.3	632,135
StoreNode01.SLHOST2	DistributedCache	☑	MACHINE-SAFE	☑	0.0	1,223,758
StoreNode01.SLHOST4	DistributedCache	☑	MACHINE-SAFE	☑	0.3	3,676,489
StoreNode04.SLHOST1	DistributedCache	☑	MACHINE-SAFE	☑	0.1	640,945
StoreNode04.SLHOST4	DistributedCache	☑	MACHINE-SAFE	☑	0.5	3,647,694

### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

- Cluster** Select a cluster to display.
- Summary** Click to view the "Single Service Summary" display.
- Service** Select a service to display.
- Host** Select a host.
- Class** Select the type of node to display: All, Storage or Process nodes.

**Cache Service Detail by Node:**

The columns in this table, with the exception of **Location**, come from Service and Node MBeans.

**Location** is a unique identifier for each node and defined as: **member\_name.machine.rack.site**.

For details on attributes of these MBeans go to: [http://download.oracle.com/otn\\_hosted\\_doc/coherence/350/com/tangosol/net/management/Registry.html](http://download.oracle.com/otn_hosted_doc/coherence/350/com/tangosol/net/management/Registry.html).

<b>Location</b>	A unique identifier for each node. It is defined as: <b>member_name.machine.rack.site</b> .
<b>Service</b>	The name of the service.
<b>Running</b>	Indicates that the service is running when checked.
<b>Metric</b>	The high availability status of the service: <ul style="list-style-type: none"> <li>● <b>ENDANGERED</b>: There is potential data loss in the cluster if a node goes offline.</li> <li>● <b>NODE-SAFE</b>: There is no risk of data loss in the cluster if a node goes offline (or is taken offline using kill-9). The data is replicated across multiple nodes and remains available in the cluster.</li> <li>● <b>MACHINE-SAFE</b>: There is no risk of data loss in the cluster if a machine goes offline (or is taken offline using kill-9). The data is replicated across multiple machines and remains available in the cluster.</li> <li>● <b>RACK-SAFE</b>: There is no risk of data loss in the cluster if a rack goes offline (or is taken offline using kill-9). The data is replicated across multiple racks and remains available in the cluster.</li> <li>● <b>SITE-SAFE</b>: There is no risk of data loss in the cluster if a site goes offline (or is taken offline using kill-9). The data is replicated across multiple sites and remains available in the cluster.</li> </ul>
<b>Time Range</b>	Select a time range from the drop down menu varying from <b>2 Minutes</b> to <b>Last 7 Days</b> , or display <b>All Data</b> . To specify a time range, click Calendar <input type="text"/> .

By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

**Process Nodes:  
History Heatmap  
of Service Metric**

Color of the cells represents the relative value of the selected Metric for a given process node; a darker shade is a larger value.  
The value of the Metric is displayed over the specified History for all process nodes in the selected Service.

**Storage Nodes:  
History Heatmap  
of Service Metric**

Color of the cells represents the relative value of the selected Metric for a given storage node; a darker shade is a larger value.  
The value of the Metric is displayed over the specified History for all storage nodes in the selected Service.

## Federated Clusters

Federated Clusters displays present high-level and detailed cache performance metrics for the cluster. Performance statistics are derived from the cluster Destination and Origin MBeans. Destination information shows how efficiently each node in the local cluster participant is sending data to each destination cluster participant. Origin information shows how efficiently each node in the local cluster participant is receiving data from destination cluster participants.

Use these displays to quickly assess total utilization and throughput metrics for all caches in the cluster.

- [“Federated Destination Detail” on page 743](#): Shows current information for all participating nodes for a selected cluster.
- [“Federated Destination Summary” on page 746](#): Shows current information and trended historical rate information.
- [“Federated Origin Detail” on page 748](#): Shows current information for all participating nodes for a selected cluster.
- [“Federated Origin Summary” on page 751](#): Shows current information and trended historical rate information.

## Federated Destination Detail

Table shows performance and utilization data, such as bandwidth usage and bytes sent, for Federated Destinations on the selected cluster. Use this display to do high level utilization analysis. Each row is a different Destination MBean. Click a row to see details in the “[Federated Destination Summary](#)” display. Sort data by the highest and lowest values of interest by clicking on the column heading.

Location ▲	BytesSentSecs	ConnectRetryTimeoutMillis	Connection	CurrentBandwidth	Delta
1.SLNB50	4,401	0	Cluster1Jmx	0.0	0.0
2.SLNB50	4,615	0	Cluster1Jmx	0.0	0.0
3.SLNB50	6,210	0	Cluster1Jmx	0.0	0.0

### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

### Filter By:

- Cluster:** Select a cluster from the drop-down menu.
- Host:** Select a host from the drop-down menu.

### Federated Destination Detail by Node

- Location** A unique identifier for each node. It is defined as: **member\_name.machine.rack.site**.
- BytesSentSecs** The number of bytes sent per second.
- ConnectRetryTimeoutMillis** The configured connect retry timeout.
- Connection** The name of the JMX connection used to access the cluster data.

<b>CurrentBandwidth</b>	The current amount of bandwidth being used, in megabits per second, for sending replicate message.
<b>DeltaReplicateAllTotalTime</b>	The difference in the total amount of time the <b>replicateAll</b> request took since the last data sample.
<b>DeltaTIME_STAMP</b>	The amount of time since the last data sample.
<b>DeltaTotalBytesSent</b>	The difference in the total number of bytes sent since the last data sample.
<b>DeltaTotalEntriesSent</b>	The difference in the total number of entries sent since the last data sample.
<b>DeltaTotalErrorResponses</b>	The difference in the total number of error responses since the last data sample.
<b>DeltaTotalMsgSent</b>	The difference in the total number of messages sent since the last data sample.
<b>DeltaTotalMsgUnacked</b>	The difference in the total number of unacknowledged messages since the last data sample.
<b>DeltaTotalRecordsSent</b>	The difference in the total number of records sent since the last data sample.
<b>ErrorDescription</b>	A description of the error. A value exists only if the sender is in an error state.
<b>EstimatedReplicateAllRemainingTime</b>	The estimated remaining time, in milliseconds, to complete the <b>replicateAll</b> request.
<b>Expired</b>	When checked, this connection is expired due to inactivity.
<b>GeoIp</b>	The Geo-IP metadata
<b>HostName</b>	The name of the host.
<b>MaxBandwidth</b>	The maximum amount of bandwidth per second, in megabits, for sending replicate message, where <b>-1.0</b> means the maximum bandwidth is not specified.
<b>Member</b>	The member information of the destination node.
<b>MemberName</b>	The name of the member.
<b>MsgApplyTimePercentileMillis</b>	The 90-percentile value, in milliseconds, of the time taken to apply the replication messages on the destination.
<b>MsgNetworkRoundTripTimePercentileMilli s</b>	The 90-percentile value, in milliseconds, of the time taken by transmission of replication messages and the corresponding ack messages over the network.
<b>MsgSentSecs</b>	The number of messages sent per second.
<b>Name</b>	The sender name.
<b>ParticipantType</b>	The participant type. Valid types are <b>cluster</b> and <b>interceptor</b> .
<b>RateReplicateAllTotalTime</b>	The number of <b>replicateAll</b> requests per second.
<b>RateTotalBytesSent</b>	The total number of bytes sent per second.
<b>RateTotalEntriesSent</b>	The total number of entries sent per second.
<b>RateTotalErrorResponses</b>	The total number of error responses per second.
<b>RateTotalMsgSent</b>	The total number of messages sent per second.
<b>RateTotalMsgUnacked</b>	The total number of unacknowledged messages per second.

<b>RateTotalRecordsSent</b>	The total number of records sent per second.
<b>RecordBacklogDelayTimePercentileMillis</b>	The 90-percentile value , in milliseconds, of the time the journal records are in the cache waiting to be replicated.
<b>ReplicateAllPercentComplete</b>	The percent of work completed for a <b>replicateAll</b> request.
<b>ReplicateAllTotalTime</b>	The total amount of time the <b>replicateAll</b> request took, in milliseconds.
<b>SendTimeoutMillis</b>	The configured send timeout.
<b>State</b>	The participant state, where: <b>0</b> is Ok <b>1</b> is Warning <b>2</b> is Error
<b>Status</b>	The participant status.
<b>TIME_STAMP</b>	The date and time of the data update.
<b>TotalBytesSent</b>	The total number of bytes sent.
<b>TotalEntriesSent</b>	The total number of cache entries sent.
<b>TotalErrorResponses</b>	The total number of responses with an error.
<b>TotalMsgSent</b>	The total number of replication messages sent. A replication message might contain multiple journal records
<b>TotalMsgUnacked</b>	The total number of unacknowledged replication messages.
<b>TotalRecordsSent</b>	The total number of journal records sent. A journal record might consist of multiple cache entries that are part of the same transaction.
<b>name</b>	The destination cluster name.
<b>nodeid</b>	The unique identifier for the node.
<b>service</b>	The Federated Service name.
<b>subType</b>	The Federated Service sub-type.
<b>type</b>	The Coherence MBean type (Federation, in this case).

## Federated Destination Summary

Detailed performance and utilization data, such as bandwidth usage and bytes sent per second, for a Federated Destinations location. Use this display to do low level utilization analysis. Check the metrics for to determine whether more capacity is needed.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

- Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

### Filter By:

- Cluster:** Select a cluster from the drop-down menu.
- Host:** Select a host from the drop-down menu.
- Location:** Select a location from the drop-down menu. **Location** is a unique identifier for each node and defined as: **member\_name.machine.rack.site**.
- Id:** The unique identifier for the node.

<b>Participant Type</b>	The participant type. Valid types are <b>cluster</b> and <b>interceptor</b> .
<b>State</b>	The participant state, where: <b>0</b> is Ok <b>1</b> is Warning <b>2</b> is Error
<b>Bytes Sent Secs</b>	The number of bytes sent per second.
<b>Connect Retry Timeout (ms)</b>	The configured connect retry timeout.
<b>Current Bandwidth</b>	The current amount of bandwidth being used, in megabits per second, for sending replicate message.
<b>Estimated Replicate All Remaining Time</b>	The estimated remaining time, in milliseconds, to complete the <b>replicateAll</b> request.
<b>Geo IP</b>	The Geo-IP metadata
<b>Max Bandwidth</b>	The maximum amount of bandwidth per second, in megabits, for sending replicate message, where <b>-1.0</b> means the maximum bandwidth is not specified.
<b>Status</b>	The participant status.
<b>Name</b>	The sender name.
<b>Msg Apply Time Percentile (ms)</b>	The 90-percentile value, in milliseconds, of the time taken to apply the replication messages on the destination.
<b>Msgs Sent Secs</b>	The number of messages sent per second.
<b>Record Backlog Delay Time Percentile (ms)</b>	The 90-percentile value, in milliseconds, of the time the journal records are in the cache waiting to be replicated.
<b>Replicate All Percentile Complete</b>	The percent of work completed for a <b>replicateAll</b> request.
<b>Replicate All Total Time</b>	The total amount of time the <b>replicateAll</b> request took, in milliseconds.
<b>Send Timeout (ms)</b>	The configured send timeout.
<b>Error Description</b>	A description of the error. A value exists only if the sender is in an error state.

### Trend Graph

Select a location from the drop-down menu to populate the trend graph. **Location** is a unique identifier for each node and defined as: **member\_name.machine.rack.site**.

**RateReplicateAllTotalTime:** Traces the total number of **replicateAll** requests per second.

**RateTotalBytesSent:** Traces the total number of bytes sent per second.

**RateTotalEntriesSent:** Traces the total number of entries sent per second.

**RateTotalErrorResponses:** Traces the total number of error responses per second.

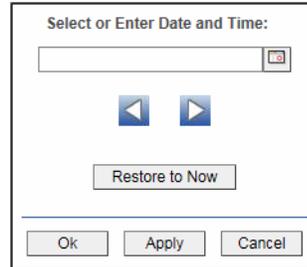
**RateTotalMsgSent:** Traces the total number of messages sent per second.

**RateTotalMsgUnacked:** Traces the total number of unacknowledged messages per second.

**RateTotalRecordsSent:** Traces the total number of records sent per second.

**ReplicateAllPercentComplete:** Traces the percent of completed **replicateAll** requests.

- Start Time** The date and time the location was started. **Location** is a unique identifier for each node and defined as: **member\_name.machine.rack.site**.
- Base at Zero** Use zero for the Y axis minimum for all graphs.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar  .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Federated Origin Detail

Table shows performance and utilization data, such as bandwidth usage and bytes sent, for Federated Origins on the selected cluster. Use this display to do high level utilization analysis. Each row is a different Origin MBean. Click a row to see details in the ["Federated Origin Summary"](#) display. Sort data by the highest and lowest values of interest by clicking on the column heading.



**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- ,  open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Filter By:**

- Cluster:** Select a cluster from the drop-down menu.
- Host:** Select a host from the drop-down menu.

**Federated Origin Detail by Node**

- Location** A unique identifier for each node. It is defined as: **member\_name.machine.rack.site.**
- BytesReceivedSecs** The number of bytes received per second.
- Connection** The name of the JMX connection used to access the cluster data.

<b>DeltaTIME_STAMP</b>	The amount of time since the last data sample.
<b>DeltaTotalBytesReceived</b>	The difference in the total number of bytes received since the last data sample.
<b>DeltaTotalEntriesReceived</b>	The difference in the total number of entries received since the last data sample.
<b>DeltaTotalMsgReceived</b>	The difference in the total number of messages received since the last data sample.
<b>DeltaTotalMsgUnacked</b>	The difference in the total number of unacknowledged messages since the last data sample.
<b>DeltaTotalRecordsReceived</b>	The difference in the total number of records received since the last data sample.
<b>Expired</b>	When checked, this connection is expired due to inactivity.
<b>HostName</b>	The name of the host.
<b>Member</b>	The member information of the destination node.
<b>MemberName</b>	The name of the member.
<b>MsgApplyTimePercentileMillis</b>	The 90-percentile value, in milliseconds, of the time taken to apply the replication messages on the origin.
<b>MsgReceivedSecs</b>	The number of messages received per second.
<b>RateReplicateAllTotalTime</b>	The number of <b>replicateAll</b> requests per second.
<b>RateTotalBytesReceived</b>	The total number of bytes received per second.
<b>RateTotalEntriesReceived</b>	The total number of entries received per second.
<b>RateTotalMsgReceived</b>	The total number of messages received per second.
<b>RateTotalMsgUnacked</b>	The total number of unacknowledged messages per second.
<b>RateTotalRecordsReceived</b>	The total number of records received per second.
<b>RecordBacklogDelayTimePercentileMillis</b>	The 90-percentile value, in milliseconds, of the time the journal records are in the cache waiting to be replicated.
<b>TIME_STAMP</b>	The date and time of the data update.
<b>TotalBytesReceived</b>	The total number of bytes received.
<b>TotalEntriesReceived</b>	The total number of cache entries received.
<b>TotalErrorResponses</b>	The total number of responses with an error.
<b>TotalMsgReceived</b>	The total number of replication messages received. A replication message might contain multiple journal records
<b>TotalMsgUnacked</b>	The total number of unacknowledged unacknowledged messages.
<b>TotalRecordsReceived</b>	The total number of journal records received. A journal record might consist of multiple cache entries that are part of the same transaction.
<b>name</b>	The destination cluster name.
<b>nodeid</b>	The unique identifier for the node.
<b>service</b>	The Federated Service name.

- subType** The Federated Service sub-type.
- type** The Coherence MBean type (Federation, in this case).

### Federated Origin Summary

Detailed performance and utilization data, such as bandwidth usage and received per second, for a Federated Origin location. Use this display to do low level utilization analysis. Check the metrics for to determine whether more capacity is needed.



**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

**Filter By:**

The display might include these filtering options:

- Cluster:** Select a cluster from the drop-down menu.

<b>Host:</b>	Select a host from the drop-down menu.
<b>Location:</b>	Select a location from the drop-down menu. <b>Location</b> is a unique identifier for each node and defined as: <b>member_name.machine.rack.site</b> .
<b>Bytes Received Secs</b>	The number of bytes received per second.
<b>Msg Apply Time Percentile (ms)</b>	The 90-percentile value, in milliseconds, of the time taken to apply the replication messages on the origin.
<b>Msgs Received Secs</b>	The number of messages received per second.
<b>Record Backlog Delay Time Percentile (ms)</b>	The 90-percentile value, in milliseconds, of the time the journal records are in the cache waiting to be replicated.
<b>Total Bytes Received</b>	The total number of bytes received.
<b>Total Entries Received</b>	The total number of cache entries received.
<b>Total Msg Received</b>	The total number of replication messages received. A replication message might contain multiple journal records.
<b>Total Msg Unacked</b>	The total number of unacknowledged replication messages.
<b>Total Records Received</b>	The total number of journal records received. A journal record might consist of multiple cache entries that are part of the same transaction.

**Trend Graph**

Select a location from the drop-down menu to populate the trend graph. **Location** is a unique identifier for each node and defined as: **member\_name.machine.rack.site**.

**RateReplicateAllTotalTime:** Traces the total number of **replicateAll** requests per second.

**RateTotalBytesReceived:** Traces the total number of bytes received per second.

**RateTotalEntriesReceived:** Traces the total number of entries received per second.

**RateTotalErrorResponses:** Traces the total number of error responses per second.

**RateTotalMsgReceived:** Traces the total number of messages received per second.

**RateTotalMsgUnacked:** Traces the total number of unacknowledged messages per second.

**RateTotalRecordsReceived:** Traces the total number of records received per second.

**ReplicateAllPercentComplete:** Traces the percent of completed **replicateAll** requests.

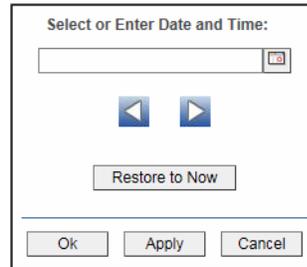
**Start Time**            The start date and time.

**Base at Zero**

Use zero for the Y axis minimum for all graphs.

**Time Range**

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

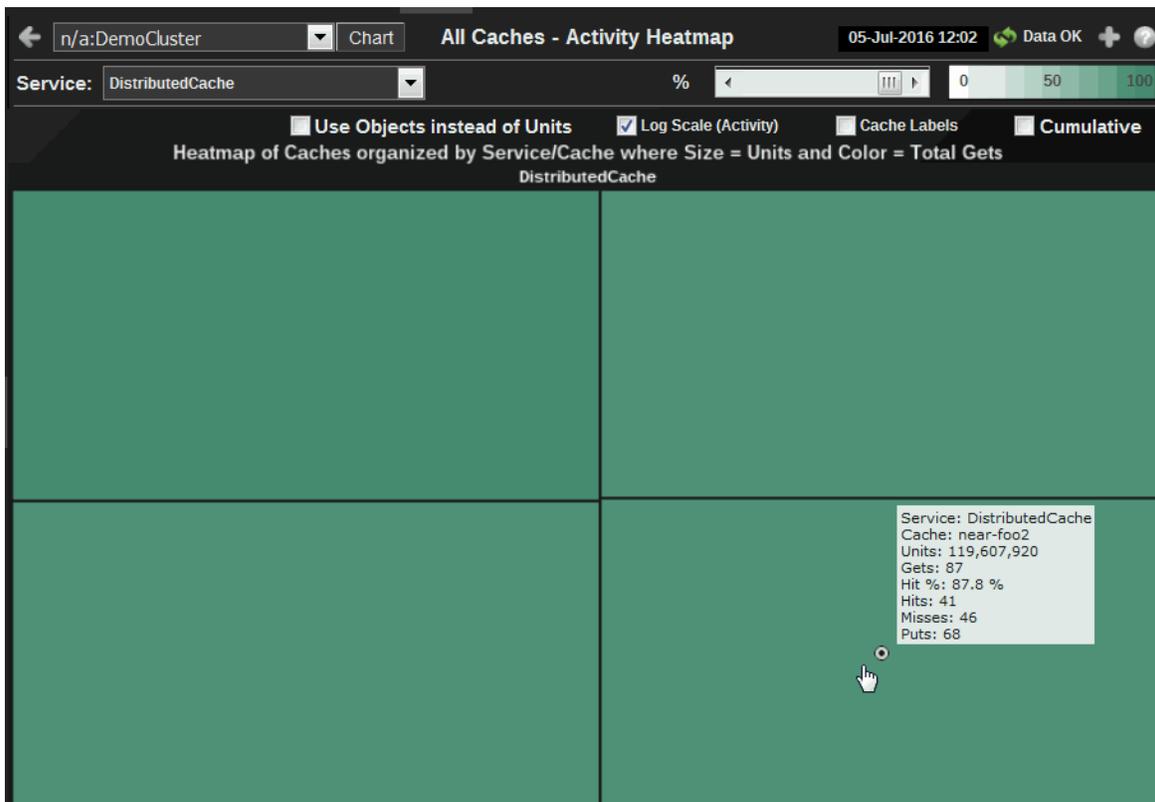
## All Caches

All Caches displays present high-level cache performance metrics for the cluster. Use the All Caches displays to quickly assess total utilization metrics for all caches in the cluster.

- ["All Caches Heatmap" on page 754](#): Heatmap of caches by service where size represents Units and color represents Total Gets%.
- ["Storage Nodes Cache Map" on page 755](#): Heatmap of memory usage on storage nodes by service where size represents Units and color represents Units Used%.
- ["Current Size Chart" on page 757](#): Bar chart/table sorted by caches with largest size displays current size/capacity metrics.
- ["Current Activity Chart" on page 758](#): Bar chart/table sorted by caches with greatest activity displays current activity metrics.

## All Caches Heatmap

Heatmap of cache size and activity organized by service: Size = Number of Units or Objects, Color = Percent of Total Gets.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

- Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

<b>Cluster</b>	Select a cluster to display.
<b>Chart</b>	Toggle between heatmap view and chart view.
<b>Service</b>	Select a service to display.
<b>%</b>	Set the activity percentage that maps to the maximum color value. Percentages greater than this value map to the maximum color value.
<b>Use Objects Instead of Units</b>	Select to use Objects instead of Units for heatmap cell sizing and mouseover tool-tips.

**Log Scale (Activity)**

Color of the cells represents the relative value of the selected Metric for a given process node; a darker shade is a larger value.  
The value of the Metric is displayed over the specified History for all process nodes in the selected Service.

**Storage Nodes: History Heatmap of Service Metric**

Color of the cells represents the relative value of the selected Metric for a given storage node; a darker shade is a larger value.  
The value of the Metric is displayed over the specified History for all storage nodes in the selected Service.

**Cache Labels**

Select to display cache labels.

**Cumulative**

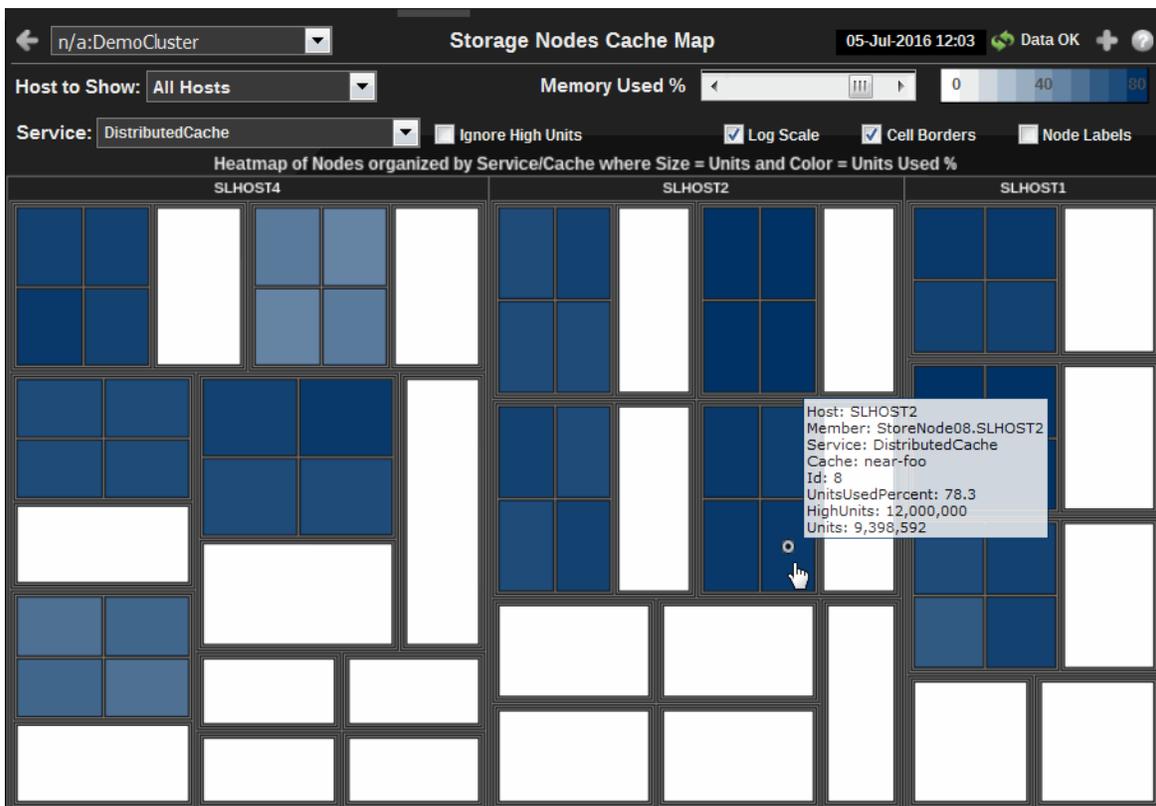
Select to show cumulative statistics for each cache.

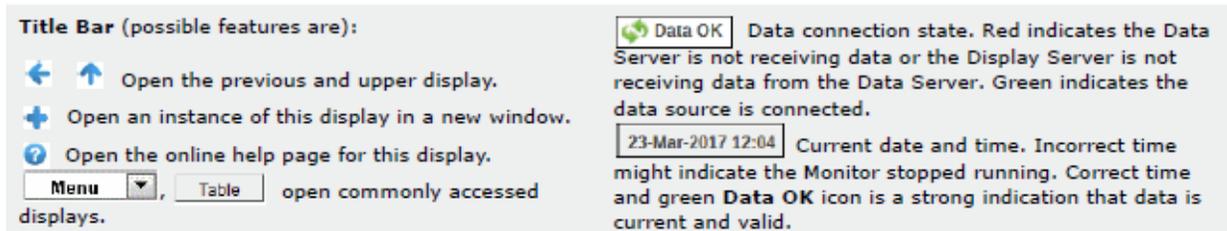
**Heatmap of Caches organized by Service/Cache**

Activity heatmap where the activity metric is TotalGets. The levels of this heatmap are Service>Cache. The size of the cells is based on Units. The size of aggregate cells is based on the sum of the Units used by its component cells. The color of the cache cells is based on TotalGets.

**Storage Nodes Cache Map**

Heatmap of memory usage on all storage nodes organized by service: Size = Number of Units, Color = Percent of Units Used.

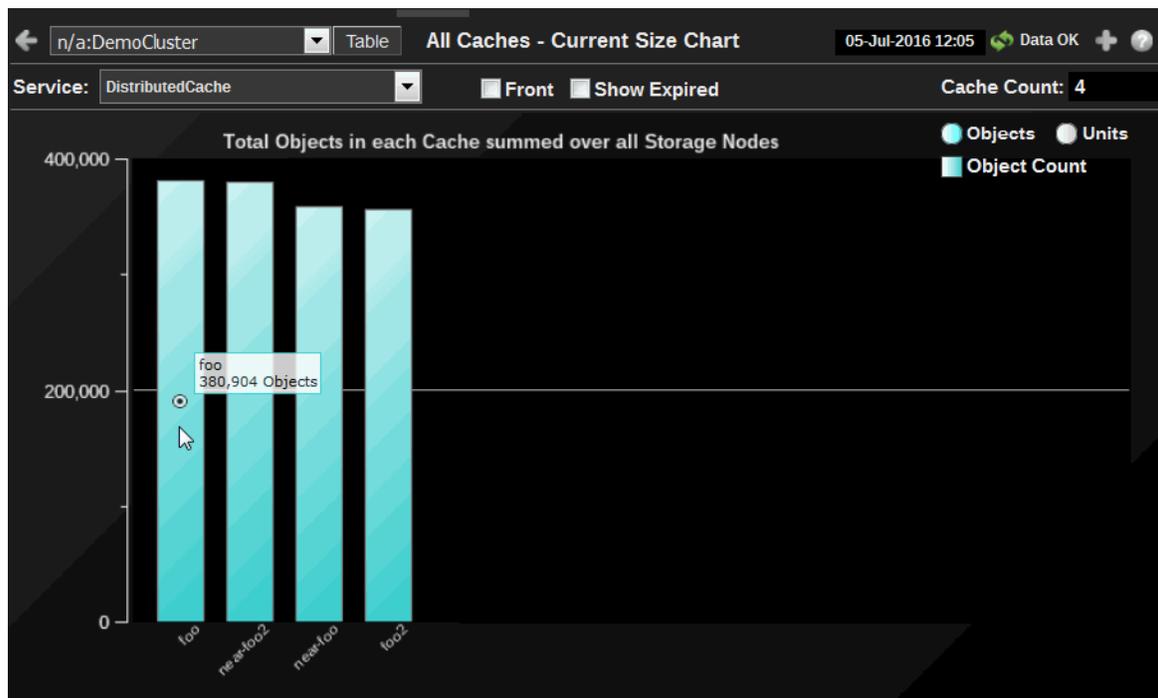




<b>Cluster</b>	Select a cluster to display.
<b>Host to Show</b>	Select a host to display.
<b>Memory Used%</b>	Set the memory used percentage that maps to the maximum color value. Percentages greater than this value map to the maximum color value.
<b>Service</b>	Select a service to display, or select All Services. NOTE: When you select a specific service, only data for nodes running that service is displayed. This enables you to view services that only run on a subset of nodes.
<b>Ignore High Units</b>	Select to remove High Units from calculations. This results in all caches having 100% units used. The color of cache cells represents units used instead of percent Units used when this option is selected.
<b>Log Scale</b>	Enable to use a logarithmic scale for the Y axis. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.
<b>Cell Borders</b>	Select to display heatmap cell borders.
<b>Node Labels</b>	Select to display node labels.
<b>Heatmap of Nodes organized by Service/Cache</b>	A heatmap of memory usage. The levels of this heatmap are <b>Host&gt;Node&gt;Service&gt;Cache</b> . The size of the cells is based on Units. The size of aggregate cells is based on the sum of the Units used by its component cells. The color of cache cells is based on the percent of Units used unless Ignore High Units is selected.

## Current Size Chart

Toggle between bar chart and table views that present the latest values of total objects and total nits for each cache in the selected service.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

<b>Cluster</b>	Select a cluster to display.
<b>Table</b>	Toggle between chart view and table view.
<b>Service</b>	Select a service to display.
<b>Front</b>	Select for front tier, deselect for back tier.
<b>Cache Count</b>	Number of caches in the selected server. This is not available in the Table view.

**Current Size Chart**

Total Objects in each Cache summed over all Storage Nodes. This is the default view. Toggle between totals for Object Count and Units Used.

Click the **Table** to view Current Size Table.

**Objects** shows the total number of objects in this cache (Object Count).

**Units** shows the highest number of units before evictions occur.

**Ignore High Units** removes High Units bars from view.

**Current Size Table**

Totals for each Cache over all Storage Nodes. Click Chart to view Current Size Chart.

**shortCacheName** Abbreviated name of cache

**tier** Front or back

**Objects** Total number of objects in this cache

**Units** Total number of units (typically bytes) in this cache

**LowUnits** Low limit for cache evictions

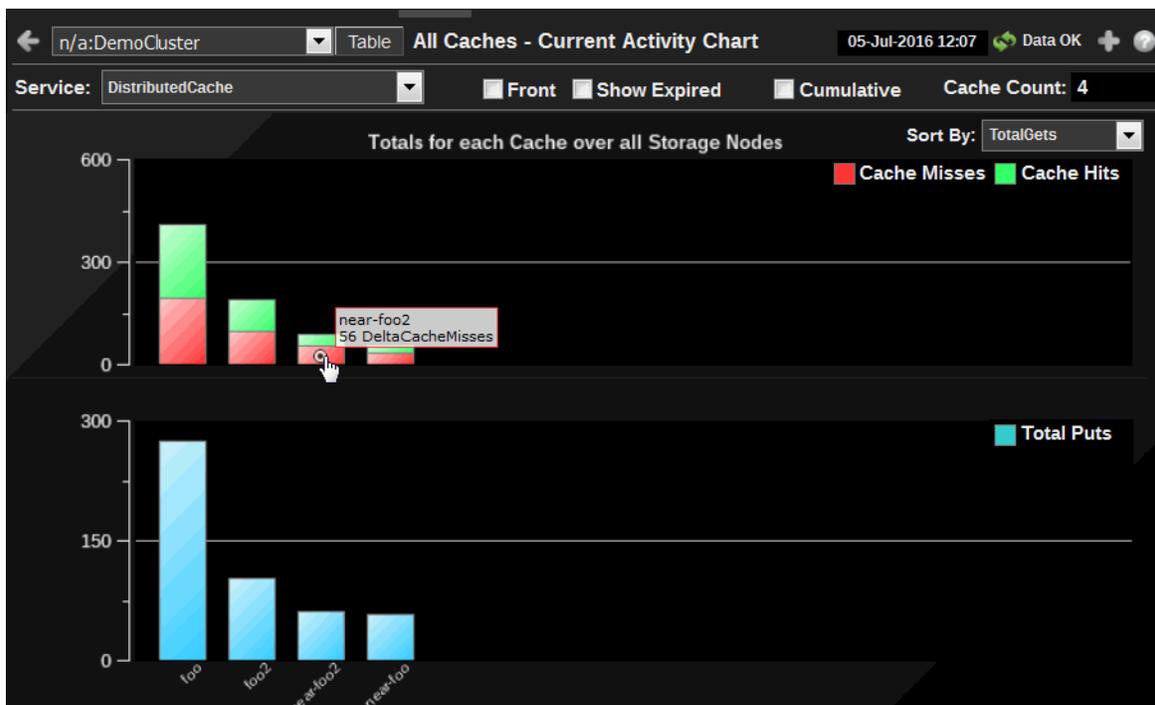
**HighUnits** Highest number of units before evictions occur

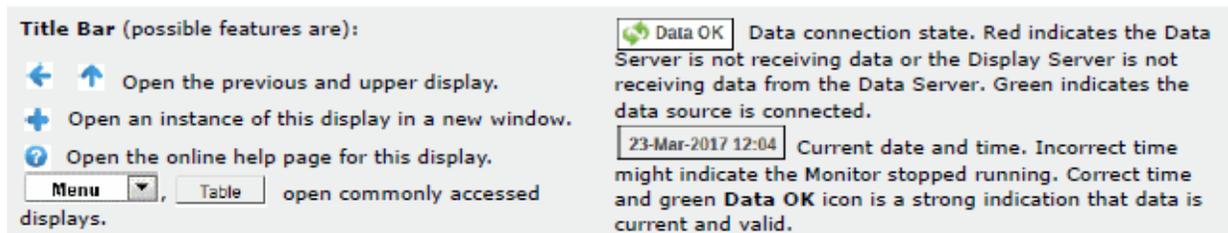
**Service Name** of selected service(s).

**Name** Full name of cache

**Current Activity Chart**

Toggle between bar chart and table views that present the latest values for activity metrics for each cache in the selected service.





<b>Cluster</b>	Select a cluster to display.
<b>Table</b>	Toggle between chart view and table view.
<b>Service</b>	Select a service to display.
<b>Front</b>	Select for front tier, deselect for back tier.
<b>Cache Count</b>	Number of caches in the selected server. This is not available in the Table view.
<b>Cumulative</b>	Select to show cumulative statistics for each node since the start of the node.
<b>Current Activity Chart</b>	Totals for Cache summed over all Storage Nodes. This is the default view. Toggle to Table view. Sort by: <b>Objects</b> shows the total number of objects in this cache (Object Count). <b>Units</b> shows the highest number of units before evictions occur. <b>Ignore High Units</b> removes High Units bars from view.
<b>Current Activity Table</b>	Totals for each Cache over all Storage Nodes. Toggle to Chart view. Sort by: <b>Cache</b> Abbreviated name of cache <b>tier</b> Front or back <b>Hits</b> Total number of successful gets <b>Misses</b> Total number of failed gets <b>Gets</b> Total requests for data from this cache <b>Puts</b> Total data stores into this cache <b>Hit%</b> Ratio of hits to gets <b>Service</b> Service Name <b>Cache Full Name</b> Full name of cache

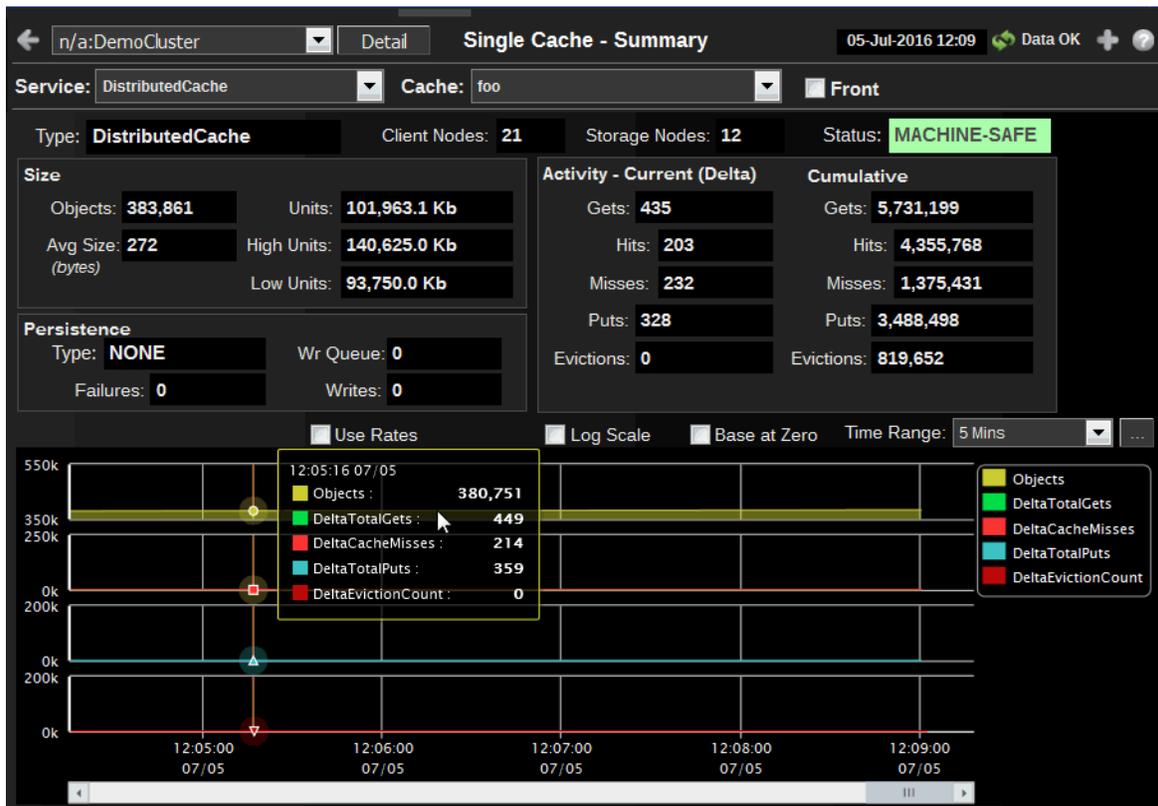
## Single Cache

Single Cache displays present detailed cache performance metrics for a single cache. Use the Single Cache displays to perform cache utilization analysis. The data in these displays can be sorted and viewed by service or cache.

- ["Single Cache Summary" on page 760](#): Perform low level utilization analysis on a single cache.
- ["Size Trends" on page 763](#): Trend chart displays size/capacity metrics.
- ["Activity Trends" on page 766](#): Trend chart displays activity metrics.
- ["Cache Detail Tables" on page 768](#): Table showing current detailed cache statistics by node.
- ["Storage Manager Detail" on page 770](#): Table showing store manager metrics.
- ["Node/Group Distribution" on page 772](#): Bar chart displays metrics showing distribution across cluster nodes or groups.
- ["Front/Back Analysis" on page 774](#): Displays metrics for the front and back tiers of a selected cache.

## Single Cache Summary

Use Single Cache - Summary display to do low level cache utilization analysis. Check the metrics for Size, Evictions and Misses to determine whether more capacity is needed. Cache Summary provides summary information about an individual cache.





<b>Cluster</b>	Select a cluster to display.
<b>Service</b>	Select a service to display.
<b>Cache</b>	Select a cache. Click the Detail button to get information specific to the selected cache.
<b>Front</b>	Select for front tier, deselect for back tier.
<b>Type</b>	The type identifier string from the ServiceMBean (ReplicatedCache, DistributedCache, etc.).
<b>Client Nodes</b>	The number of cluster nodes that do not have storage enabled.
<b>Storage Nodes</b>	Select to display storage node data in the trend graphs of this display.
<b>Type</b>	The type of cache.
<b>Storage Nodes</b>	The number of storage nodes in the cache.
<b>Status</b>	The high availability status of the service: <ul style="list-style-type: none"> <li>● <b>ENDANGERED:</b> There is potential data loss in the cluster if a node goes offline.</li> <li>● <b>NODE-SAFE:</b> There is no risk of data loss in the cluster if a node goes offline (or is taken offline using kill-9). The data is replicated across multiple nodes and remains available in the cluster.</li> <li>● <b>MACHINE-SAFE:</b> There is no risk of data loss in the cluster if a machine goes offline (or is taken offline using kill-9). The data is replicated across multiple machines and remains available in the cluster.</li> <li>● <b>RACK-SAFE:</b> There is no risk of data loss in the cluster if a rack goes offline (or is taken offline using kill-9). The data is replicated across multiple racks and remains available in the cluster.</li> <li>● <b>SITE-SAFE:</b> There is no risk of data loss in the cluster if a site goes offline (or is taken offline using kill-9). The data is replicated across multiple sites and remains available in the cluster.</li> </ul>

<b>Size</b>	<p>Units indicates memory usage for the back tier and number of objects for the front tier.</p> <p><b>Objects</b> The number of objects in the selected cache. The value is the total across all storage nodes.</p> <p><b>Avg Size</b> The average size of objects in the selected cache (in bytes if it is the back tier).</p> <p><b>Units</b> The memory usage if back tier, or number of objects if front tier. The value is the total across all storage nodes.</p> <p><b>High Units</b> Maximum memory, or number of objects allowed before Coherence starts to evict objects from the selected cache. The value is the total across all storage nodes.</p> <p><b>Low Units</b> The level of memory, or number of objects to which Coherence will reduce the cache during the eviction process. The value is the total across all storage nodes.</p>
<b>Persistence</b>	<p><b>Type</b> The persistence type for the cache. Possible values include: <b>NONE</b>, <b>READ-ONLY</b>, <b>WRITE-THROUGH</b>, and <b>WRITE-BEHIND</b>.</p> <p><b>Failures</b> The number of write (cache store) failures, including load, store and erase operations. NOTE: This value is <b>-1</b> if the persistence type is <b>NONE</b>.</p> <p><b>Wr Queue</b> The size of the queue, in kilobytes, that holds data scheduled to be written to the cache store.</p> <p><b>Writes</b> The number of objects (cache entries) written to the cache store.</p>
<b>Activity</b>	<p><b>Current:</b> Use the <b>Use Rates</b> checkbox to toggle between two value types: <b>Activity - Current (Rate)</b> and <b>Activity - Current (Delta)</b> (as labeled in the display upon selection). When the Use Rates (checkbox) is NOT selected the Delta values are shown in the Activity - Current (Delta) fields and trend graphs. Delta is the difference in the value since the last sample. When the Use Rates (checkbox) is selected the Rate values are shown in the Activity - Current (Rate) fields and trend graphs. Rate is the value per second. The Rate value is useful when the sampling time period is unknown, has changed, or has a long duration specified. For a given rate, the Rate value does not vary if the sample period changes (whereas the Delta value does vary). The Rate value enables you to directly compare rates on systems with different sample periods.</p> <p><b>Cumulative:</b> The total since the service was started for the selected cache, or since statistics were reset.</p> <p><b>Gets</b> The number of requests for data from this cache.</p> <p><b>Hits</b> The number of successful gets.</p> <p><b>Misses</b> The number of failed gets.</p> <p><b>Puts</b> The number of data stores into this cache.</p> <p><b>Evictions</b> The number of objects removed to make room for other objects.</p>
<b>Use Rates</b>	<p>Select <b>Use Rates</b> to show the Rate values in the Activity - Current (Rate) fields and trend graphs. Rate is the value per second. The Rate value is useful when the sampling time period is unknown, has changed, or has a long duration specified. For a given rate, the Rate value does not vary if the sample period changes (whereas the Delta value does vary). The Rate value enables you to directly compare rates on systems with different sample periods.</p> <p>Deselect Use Rates to show the Delta values in the <b>Activity - Current (Delta)</b> fields and trend graphs. Delta is the difference in the value since the last sample.</p>
<b>Log Scale</b>	<p>Enable to use a logarithmic scale for the Y axis. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.</p>
<b>Base at Zero</b>	<p>Use zero for the Y axis minimum for all graphs.</p>

**Time Range**

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

**Trend Graphs**

Use the Use Rates checkbox to toggle between two value types: Activity - Current (Rate) and Activity - Current (Delta) (as labeled in the display upon selection).

**Objects** The number of objects in the selected cache. The value is the total across all storage nodes.

**TotalGets** Total requests for data from this cache.

**CacheMisses** Total number of failed gets.

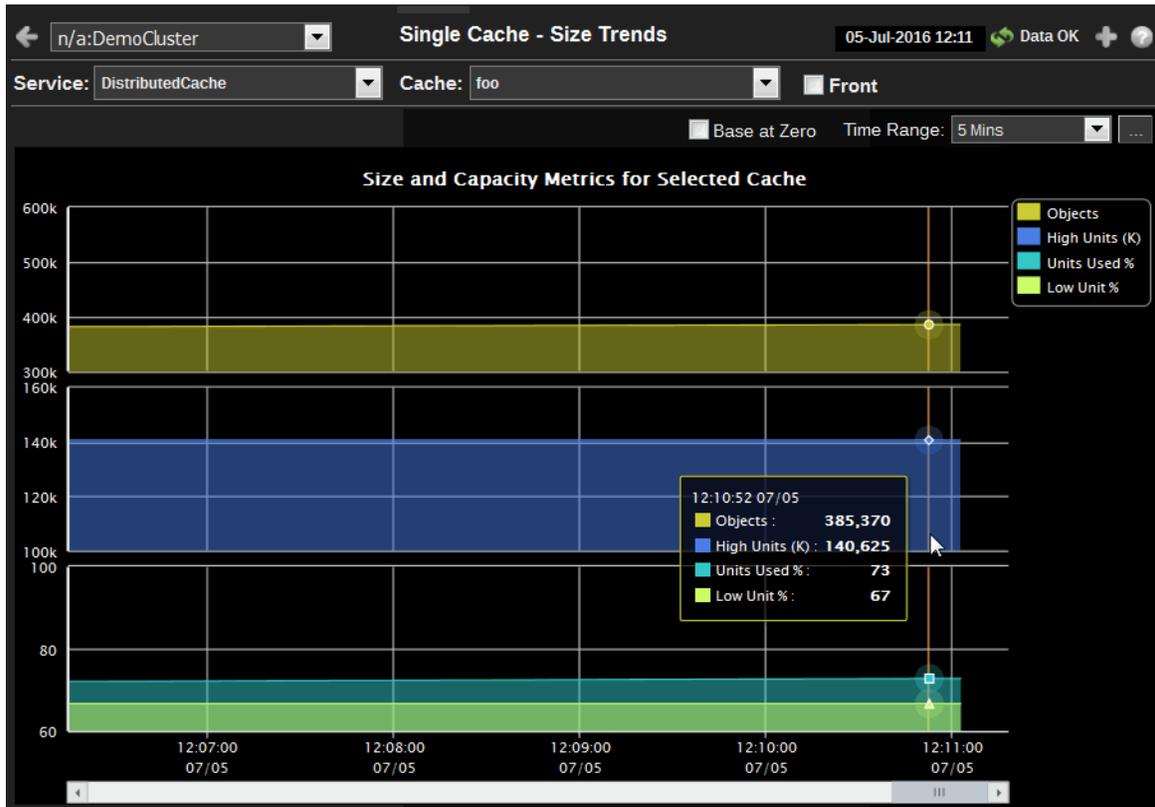
**TotalPuts** Total data stores into this cache.

**EvictionCounts** Number of objects removed from the cache to make room for other objects.

**Size Trends**

Size Trends provides a method of viewing the degree to which available cache size has been consumed. Under normal operations the cache will evict and reload objects into the cache. This will be displayed as a significant drop in the Units Used trend. However, if these drops are too frequent the application might not be performing optimally. Adding capacity and examining or modifying application usage patterns might be required. The data displayed here is a sum of all storage nodes in the cache filtered by the selected service and cache.

Try changing the High Units setting in the Cache Administration page to something like 100,000 and then see the effect on these trend charts.



**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

- Cluster** Select a cluster to display.
- Service** Select a service to display.
- Cache** Select a cache. Click the Detail button to get information specific to the selected cache.
- Front** Select for front tier, deselect for back tier.

**Base at Zero**

Use zero for the Y axis minimum for all graphs.

**Time Range**

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Activity Trends

Activity Trends provides a set of trend graphs that show the magnitude of the cache usage and the effectiveness of the implementation. If the overall effectiveness is not as desired, increasing capacity, preloading the cache and increasing the eviction time may result in improvements in cache hits. The data displayed here is a sum of all storage nodes in the cache filtered by the selected service and cache.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu Table open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

- Cluster** Select a cluster to display.
- Service** Select a service to display.
- Cache** Select a cache. Click the Detail button to get information specific to the selected cache.
- Front** Select for front tier, deselect for back tier.

**Log Scale**

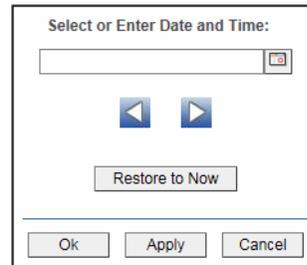
Enable to use a logarithmic scale for the Y axis. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.

**Base at Zero**

Use zero for the Y axis minimum for all graphs.

**Time Range**

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

**Activity Metrics for Selected Cache**

**Hits** The number of successful gets from this cache.

**Total Gets** Requests for data from this cache.

**Cache Misses** The number of failed gets by this cache.

**Total Puts** The number of data stores into this cache.

**Evictions** The number of objects removed from the cache to make room for other objects.

**Write Queue** The size of the queue, in kilobytes, that holds data scheduled to be written to the cache store.

## Cache Detail Tables

This display presents detailed information about the contribution that each storage node makes to the cache. Select a node in the Statistics By Node for Selected Cache table to drill down to the "Node Summary" display for that node. The data displayed here is broken down for each storage nodes in the cache filtered by the selected service and cache.

Objects	Units	LowUnits	HighUnits	Hits	Delta	Misses	Delta	Gets	Delta	Puts	Delta
389,542	105,955,424	96,000,000	144,000,000	4,366,039	282	1,383,079	193	5,749,118	475	3,500,956	310

AvgSize	Units	LowUnits	HighUnits	Hits	Delta	Misses	Delta	Gets	Delta	Puts	Delta
272	8,829,618	8,000,000	12,000,000	363,836	23	115,256	16	479,093	39	291,746	25

Location	tier	Objects	AvgSize	Units	LowUnits	HighUnits	Hits
StoreNode01.SLHOST1	back	34,541	272	9,395,152	8,000,000	12,000,000	1,669
StoreNode01.SLHOST2	back	37,492	272	10,197,824	8,000,000	12,000,000	183,778
StoreNode01.SLHOST4	back	21,943	272	5,968,496	8,000,000	12,000,000	734,269
StoreNode04.SLHOST1	back	33,379	272	9,079,088	8,000,000	12,000,000	1,628
StoreNode04.SLHOST4	back	31,346	272	8,526,112	8,000,000	12,000,000	658,192
StoreNode05.SLHOST2	back	32,555	272	8,854,960	8,000,000	12,000,000	202,794
StoreNode05.SLHOST4	back	32,889	272	8,945,808	8,000,000	12,000,000	770,567
StoreNode05n.SLHOST1	back	39,918	272	10,857,696	8,000,000	12,000,000	1,650
StoreNode05n.SLHOST2	back	33,736	272	9,176,192	8,000,000	12,000,000	200,245
StoreNode05n.SLHOST4	back	31,224	272	8,492,928	8,000,000	12,000,000	764,372
StoreNode08.SLHOST2	back	35,390	272	9,626,080	8,000,000	12,000,000	151,865
StoreNode08.SLHOST4	back	25,129	272	6,835,088	8,000,000	12,000,000	695,010

### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

- Cluster** Select a cluster to display.
- Summary** Toggle between this display and Single Cache - Summary display.
- Service** Select a service to display.
- Cache** Select a cache. Click the Detail button to get information specific to the selected cache.
- Front** Select for front tier, deselect for back tier.

**Totals for Selected Cache**

**Objects** Number of objects in this cache.  
**Units** Total number of units (typically bytes) in this cache.  
**LowUnits** Low limit for cache evictions.  
**HighUnits** Highest number of units before evictions occur.  
**Hits** Total number of successful gets.  
**Misses** Total number of failed gets.  
**Gets** Total requests for data from this cache.  
**Puts** Total data stores into this cache.

**Average for Selected Cache**

**Objects** Number of objects in this cache.  
**AvgSize** Average size of objects in this cache.  
**Units** Average number of units (typically bytes) in this cache.  
**LowUnits** Low limit for cache evictions.  
**HighUnits** Highest number of units before evictions occur.  
**Hits** Average number of successful gets.  
**Misses** Average number of failed gets.  
**Gets** Average requests for data from this cache.  
**Puts** Average data stores into this cache.

**Statistics By Node for Selected Cache**

The columns in this table, with the exception of **Location**, come from Cache and Node MBeans. **Location** is a unique identifier for each node and defined as: **member\_name.machine.rack.site**.

For details about attributes of these MBeans go to: [http://download.oracle.com/otn\\_hosted\\_doc/coherence/350/com/tangosol/net/management/Registry.html](http://download.oracle.com/otn_hosted_doc/coherence/350/com/tangosol/net/management/Registry.html).

## Storage Manager Detail

This display presents detailed information about the Storage Manager. The data displayed here is queried from the Coherence StorageManagerMBean, filtered by the selected service and cache. Click on a row in the table to open the "Storage IndexInfo View" window.

The screenshot shows the 'Storage Manager Detail' window for a cluster named 'n/a:DemoCluster'. The service is 'DistributedCache' and the cache is 'foo'. The table below shows metrics for 15 different nodes. A mouse cursor is hovering over the 'StoreNode04.SLHOST4' row.

Location	EventsDispatched	EvictionCount	InsertCount	ListenerFilterCount	ListenerKeyCount
StoreNode01.SLHOST1	0	0	34,620	0	0
StoreNode01.SLHOST2	0	38,015	75,579	0	0
StoreNode01.SLHOST4	0	146,554	168,601	0	0
StoreNode04.SLHOST1	0	0	33,472	0	0
StoreNode04.SLHOST4	0	120,540	151,971	0	0
StoreNode05.SLHOST2	0	60,349	92,983	0	0
StoreNode05.SLHOST4	0	115,337	148,315	0	0
StoreNode05n.SLHOST1	0	0	40,006	0	0
StoreNode05n.SLHOST2	0	43,208	77,031	0	0
StoreNode05n.SLHOST4	0	130,074	161,389	0	0
StoreNode08.SLHOST2	0	44,379	79,849	0	0
StoreNode08.SLHOST4	0	121,196	146,401	0	0

### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

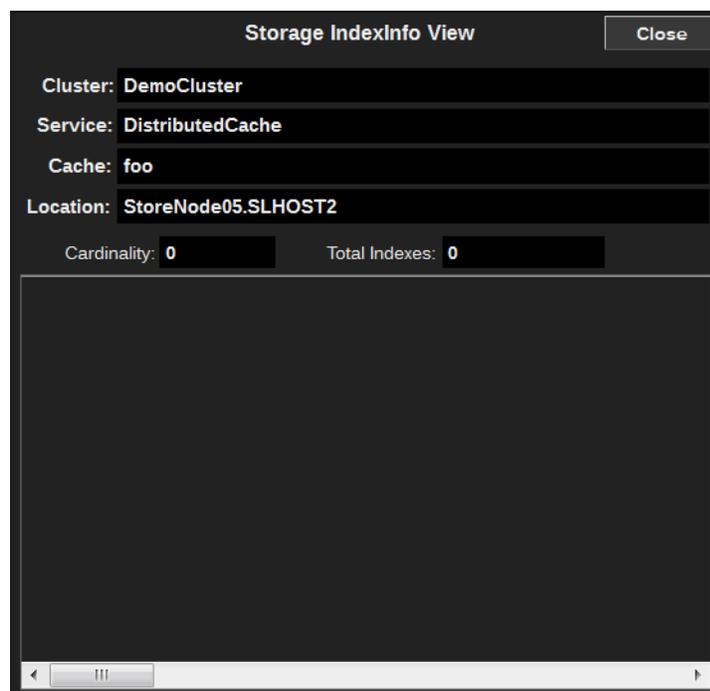
**Cluster** Select a cluster to display.

**Service** Select a service to display.

<b>Cache</b>	Select a cache. Click the Detail button to get information specific to the selected cache.
<b>Storage Manager Data</b>	<p><b>Location</b> A unique identifier for each node. It is defined as <b>member_name.machine.rack.site</b>.</p> <p><b>EventsDispatched</b> The total number of events dispatched by the Storage Manager since the last time the statistics were reset.</p> <p><b>EvictionCount</b> The number of evictions from the backing map managed by this Storage Manager caused by entries expiry or insert operations that would make the underlying backing map to reach its configured size limit.</p> <p><b>InsertCount</b> The number of inserts into the backing map managed by this Storage Manager. In addition to standard inserts caused by put and invoke operations or synthetic inserts caused by get operations with read-through backing map topology, this counter is incremented when distribution transfers move resources into the underlying backing map and is decremented when distribution transfers move data out.</p> <p><b>ListenerFilterCount</b> The number of filter-based listeners currently registered with the Storage Manager.</p> <p><b>ListenerKeyCount</b> The number of key-based listeners currently registered with the Storage Manager.</p> <p><b>ListenerRegistrations</b> The total number of listener registration requests processed by the Storage Manager since the last time the statistics were reset.</p> <p><b>LocksGranted</b> The number of locks currently granted for the portion of the partitioned cache managed by the Storage Manager.</p> <p><b>LocksPending</b> The number of pending lock requests for the portion of the partitioned cache managed by the Storage Manager.</p> <p><b>RemoveCount</b> The number of removes from the backing map managed by this Storage Manager caused by operations such as clear, remove or invoke.</p>

### Storage IndexInfo View

Click on a row in the Storage Manager Data table to open the Storage IndexInfo View window.

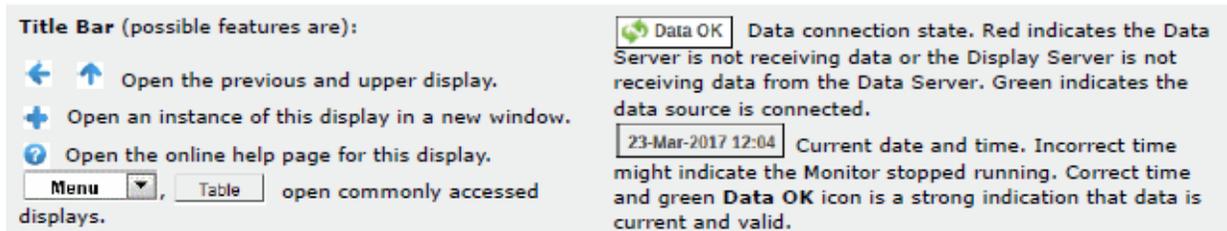


- Service**                    The name of the service.
- Cache**                     The name of the cache.
- Location Manager Data**    The location of the node associated with the cache. **Location** is a unique identifier for each node and defined as: **member\_name.machine.rack.site**.
- (Index Table)**             Each row in the table represents a unique index, where:
  - Extractor** = the index name.
  - Ordered** = true/false to indicate whether or not the data is sorted (false means the data is not sorted).
  - Size** = the number of entries in that cache whose value matches that extractor.

### Node/Group Distribution

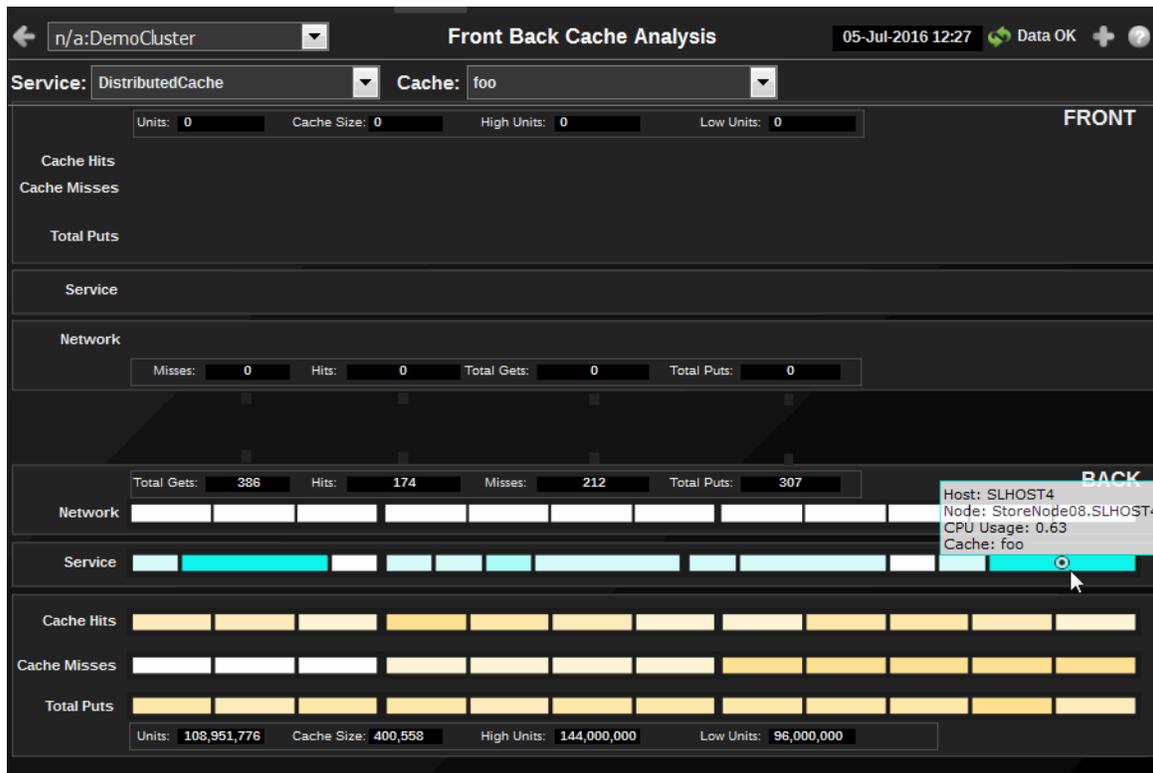
This display presents the distribution of cache activity across all storage nodes in the cluster. The buttons on the left may be used to select the metric by which all six bar charts are to be sorted. Note that the Gets, Hits, Misses, and Puts are shown in the same color as on the other Cache Analysis displays. The data displayed here is broken down for each storage nodes in the cache filtered by the selected service and cache.





<b>Cluster</b>	Select a cluster to display.
<b>Service</b>	Select a service to display.
<b>Cache</b>	Select a cache. Click the Detail button to get information specific to the selected cache.
<b>Group By</b>	Select the node group by which the data are totaled. <b>Location</b> A unique identifier for each node, defined as <b>member_name.machine.rack.site</b> . This is the default setting. <b>Gets</b> Requests for data from this cache. <b>Hits</b> Number of successful gets. <b>Misses</b> Number of failed gets. <b>Puts</b> Data stores into this cache. <b>Mem%</b> Calculated percent of memory used divided by total memory. <b>K Units</b> Units in thousand bytes.

## Front/Back Analysis



**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

- Cluster** Select a cluster to display.
- Service** Select a service to display.
- Cache** Select a cache. Click the Detail button to get information specific to the selected cache.

**FRONT/BACK****Units:**

**Front** Number of objects. The value is the total across all storage nodes for the given tier.

**Back** Memory usage. The value is the total across all storage nodes for the given tier.

**Cache Size:**

Total number of objects in the cache for the given tier (Front or Back). NOTE: Same value as Units for Front tier.

**High Units:**

**Front** Number of objects allowed before Coherence starts to evict objects from the selected cache. The value is the total across all storage nodes for the given tier.

**Back** Maximum memory allowed before Coherence starts to evict objects from the selected cache. The value is the total across all storage nodes for the given tier.

**Low Units:**

**Front** Number of objects to which Coherence will reduce the cache during the eviction process. The value is the total across all storage nodes for the given tier.

**Back** The level of memory to which Coherence will reduce the cache during the eviction process. The value is the total across all storage nodes for the given tier.

<b>Cache Hits</b>	Number of successful gets
<b>Cache Misses</b>	Number of failed gets
<b>Total Puts</b>	Data stores into this cache
<b>Service</b>	CPU usage (%) for the node.
<b>Network</b>	<b>Front</b> Sent Packet Failure Rate (%) for the node. <b>Back</b> Received Packet Failure Rate (%) for the node.
<b>Misses</b>	Number of failed gets.
<b>Hits</b>	Number of successful gets.
<b>Total Gets</b>	Total requests for data from this cache.
<b>Total Puts</b>	Total data stores into this cache.

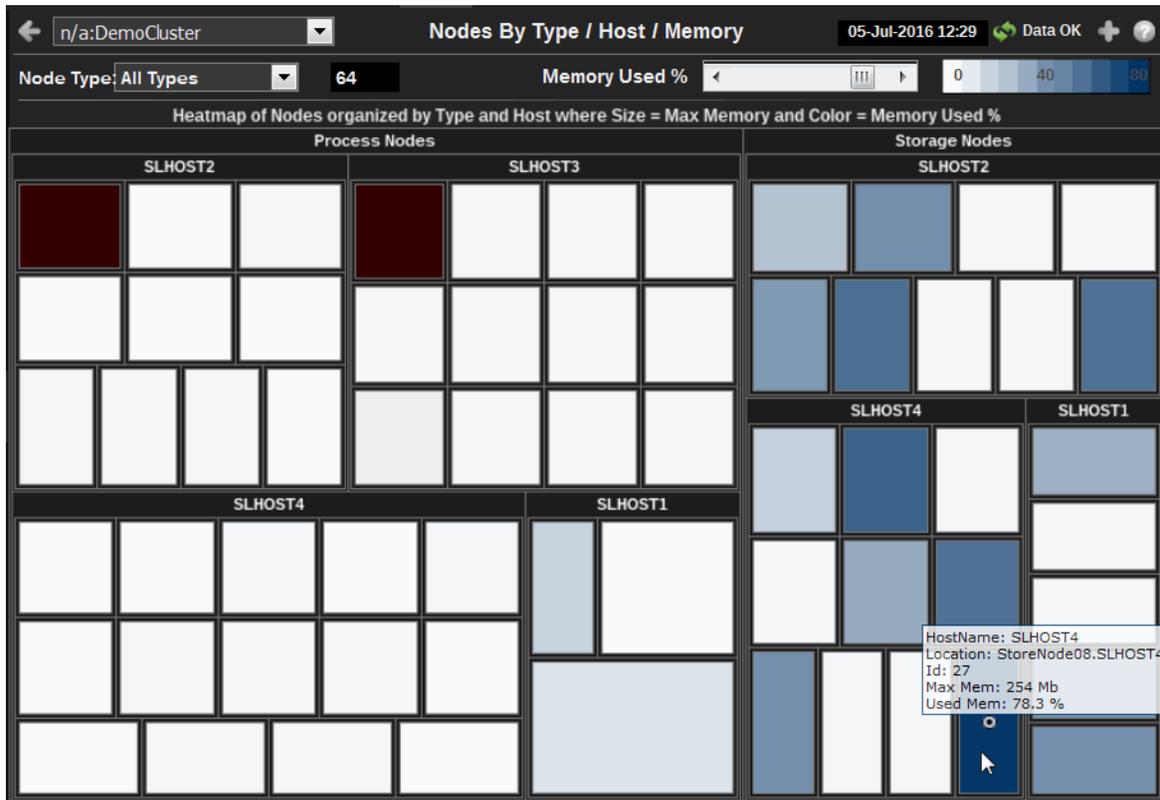
**All Nodes**

All Nodes displays present high-level node performance metrics for the cluster. Use the All Nodes displays to quickly assess total utilization metrics for all nodes in the cluster.

- ["All Nodes by Type/Host/Memory" on page 776](#): Heatmap of caches by service where size represents Max Memory and color represents percent of Memory Used.
- ["All Nodes CPU" on page 777](#): Heatmap shows CPU utilization for all nodes in the cluster.
- ["All Nodes Grid View" on page 778](#): Grid view showing information about all nodes.
- ["Communication Issues" on page 779](#): Bar chart displays current communication issues for all nodes.
- ["All Nodes - Detail" on page 781](#): Table shows current detailed statistics for all nodes.
- ["Invocation Service Detail" on page 783](#): Table shows invocation service detail for all nodes.

## All Nodes by Type/Host/Memory

Heatmap of nodes organized by Type and Host: Size = Max Memory, Color = Percent of Memory Used.



**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

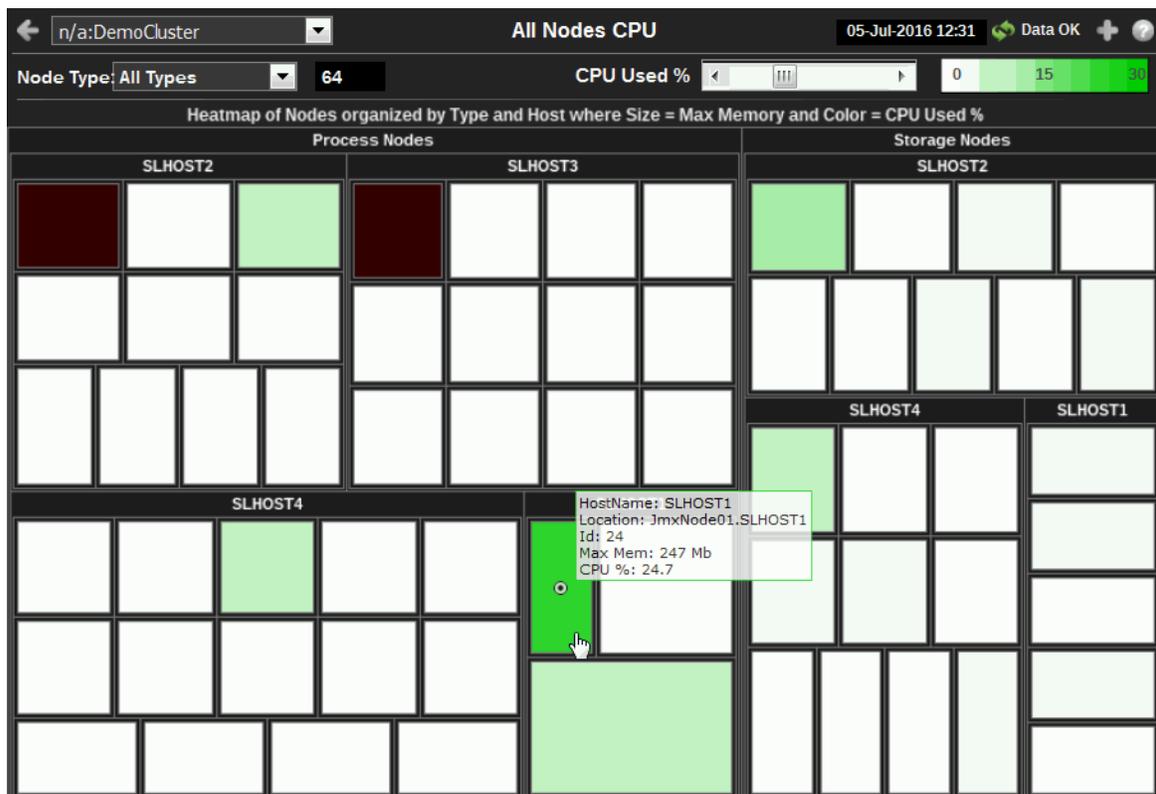
**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

- Cluster** Select a cluster to display.
- Nodes Type** Select the type of node to display: Storage Nodes, Process Nodes or All Types.
- Memory Used%** Set the memory used percentage that maps to the maximum color value. Percentages greater than this value map to the maximum color value.
- Heatmap of Nodes organized by Type/Host** A heatmap of memory usage per host.

## All Nodes CPU

Heatmap shows CPU utilization for all nodes in the cluster organized by Type and Host: Size = Max Memory, Color = Percent of CPU Used.



**Title Bar** (possible features are):

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu, Table open commonly accessed displays.

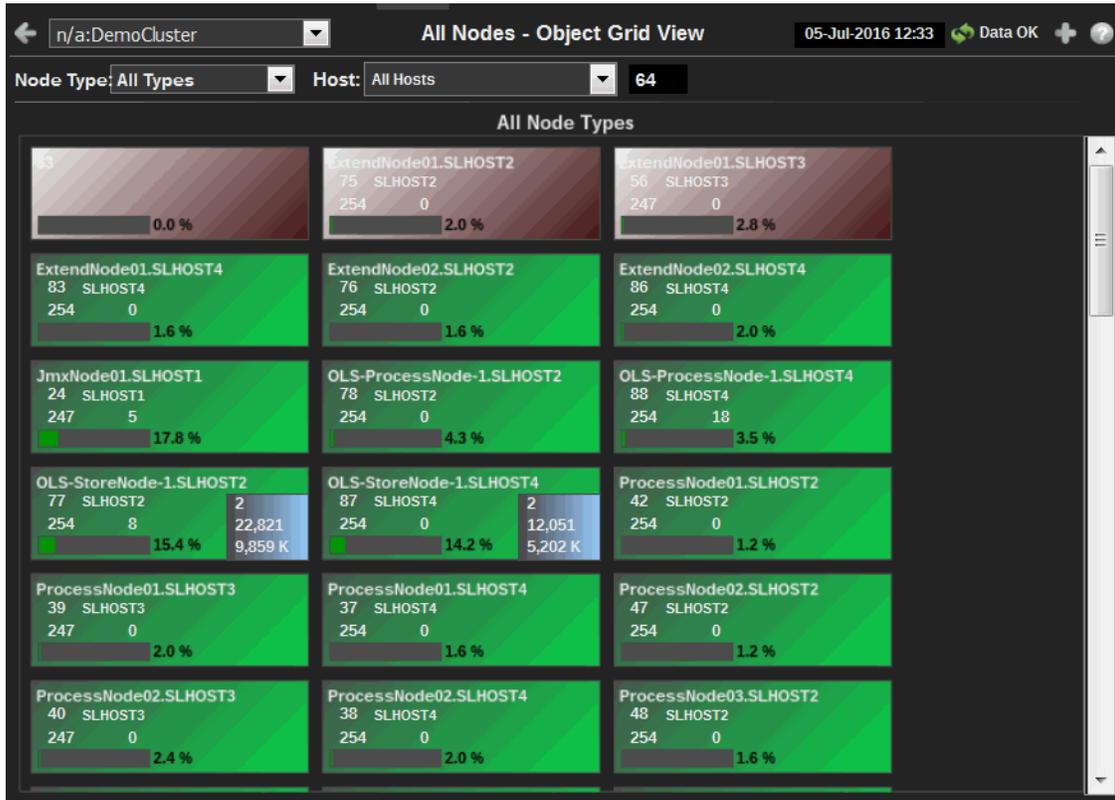
Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

- Cluster** Select a cluster to display.
- Node Type** Select the type of node to display: Storage Nodes, Process Nodes or All Types.
- CPU Used%** Set the CPU used percentage that maps to the maximum color value. Percentages greater than this value map to the maximum color value.
- Heatmap of Nodes organized by Type/Host** A heatmap of CPU usage per host.

## All Nodes Grid View

This display shows a grid view of all of the nodes in the selected Node Type.



**Title Bar (possible features are):**

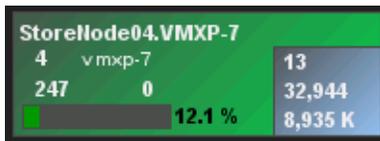
- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

- Cluster** Select a cluster to display.
- Node Type** Select the type of node to display: Storage Nodes, Process Nodes or All Types.
- Host** Select a host to display.
- Heatmap of Nodes organized by Type/Host** A heatmap of CPU usage per host.

The following icon is shown for each node in the cluster:



The icon describes the node:

- Location (**StoreNode04.VMXP-7**) A unique identifier for each node. It is defined as: **member\_name.machine.rack.site**.
- Id (**4**)
- Host name or IP (**vmpx-7**)
- Max megabytes (**247**)
- Messages queued (**0**)
- Meter and label indicating the percent of memory utilization(**12.1%**)

For storage nodes, the following are also shown (in the lower right portion of the icon):

- Number of supported caches (**13**),
- Number of objects (**32,944**)
- Amount of memory used (**8,935 K**).

## Communication Issues

This display presents detail information about communication issues by node or group. Both bar charts show the same data as the Packet Detail table. Click on a bar in either chart to drill down to the "Node Summary" display for that node.





<b>Cluster</b>	Select a cluster to display.
<b>Node Count</b>	Number of nodes in the cluster.
<b>Packets Repeated Recently</b>	Total number of repeated packets since the last update. The update rate is set by the Reporting Period.
<b>Resent Recently</b>	Total number or resent packets since the last update. The update rate is set by the Reporting Period.
<b>Reporting Period</b>	Select period varying from 30 Seconds to Last 7 Days, or display All Data.
<b>Sort By</b>	Select Packets Sent, Packets Received, Sent Failure Rate or Received Failure Rate.
<b>Sent Failure Rate/Received Failure Rate by Node/Group</b>	Packets failed to be sent by each node. Packets failed to be received by each node.
<b>Packets Sent/Received by Node/Group</b>	Packets sent by each node. Packets received by each node.

## All Nodes - Detail

This display presents detailed information about each node. This display includes information from the Coherence ClusterNodeMBean for both storage and processing nodes. Select a node in the All Node Data table to drill down to the "Node Summary" display for that node.

Location	Id	Avail MB	Max MB	Pkts Sent	Delta	Pkts Rcvd	Delta	Pkts
StoreNode08.SLHOST4	27	56	254	20,680,021	1,881	19,769,956	1,759	13
StoreNode08.SLHOST2	8	112	254	7,207,444	1,755	6,920,285	1,629	4
StoreNode07.SLHOST4	16	248	254	8,479,585	814	7,596,866	703	4
StoreNode07.SLHOST2	7	249	254	2,954,996	811	2,652,154	693	4
StoreNode06.SLHOST4	15	245	254	14,413,751	1,289	13,550,004	1,187	4
StoreNode06.SLHOST2	6	246	254	4,998,465	1,519	4,697,421	1,386	4
StoreNode05n.SLHOST4	13	144	254	9,240,557	832	8,292,349	717	2
StoreNode05n.SLHOST2	4	106	254	3,215,830	818	2,917,402	698	2
StoreNode05n.SLHOST1	17	145	247	145,793	890	146,745	749	4
StoreNode05.SLHOST4	14	112	254	9,045,804	734	8,096,728	619	6
StoreNode05.SLHOST2	5	149	254	3,130,043	799	2,826,880	694	3
StoreNode04.SLHOST4	12	163	254	20,767,959	1,825	19,859,593	1,720	6
StoreNode04.SLHOST1	19	141	247	351,071	1,934	350,140	1,798	4
StoreNode03.SLHOST4	11	249	254	8,452,999	915	7,564,521	800	4
StoreNode03.SLHOST2	3	249	254	2,911,819	901	2,606,105	784	4
StoreNode03.SLHOST1	22	236	247	151,367	890	132,011	751	4
StoreNode02.SLHOST4	10	247	254	14,330,435	1,265	13,452,552	1,148	4
StoreNode02.SLHOST2	2	245	254	4,873,858	1,357	4,569,689	1,226	4
StoreNode02.SLHOST1	23	234	247	243,086	1,248	223,462	1,109	4
StoreNode01.SLHOST4	9	97	254	9,003,123	710	8,202,539	600	5
StoreNode01.SLHOST2	1	139	254	3,074,838	682	2,822,646	572	4
StoreNode01.SLHOST1	18	169	247	145,132	848	147,525	718	4

### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

- Cluster** Select a cluster to display.
- Node Type** Select the type of nodes for which to display data: Storage Nodes, Process Nodes or All Types.
- Host** Select the host for which to display data, or select All Hosts.
- Node Count** Number of nodes for which data is currently displayed.
- Total Avail MB** Total available memory of all nodes in the cluster.
- Max** Total max memory of all nodes in the cluster.

### All Node Types (MBean Detail Data)

- **Location** A unique identifier for each node. It is defined as: **member\_name.machine.rack.site**.
- **Id** The short member id that uniquely identifies this member.
- **Avail MB** The amount of available memory for this node in MB.

- **Max MB** The maximum amount of memory for this node in MB.
- **Pkts Sent** The cumulative number of packets sent by this node since the node statistics were last reset.
- **Delta** The number of packets sent by this node since the last update.
- **Pkts Rcvd** The cumulative number of packets received by this node since the node statistics were last reset.
- **Delta** The number of packets received by this node since the last update.
- **Pkts Rptd** The cumulative number of duplicate packets received by this node since the node statistics were last reset.
- **Delta** The number of duplicate packets received by this node since the last update.
- **Pkts Resent** The cumulative number of packets resent by this node since the node statistics were last reset.
- **Delta** The number of packets resent by this node since the last update.
- **Timestamp** The date and time (in cluster time) that this member joined the cluster.
- **Pub Succ Rate** The publisher success rate for this node since the node statistics were last reset. Publisher success rate is a ratio of the number of packets successfully delivered in a first attempt to the total number of sent packets. A failure count is incremented when there is no ACK received within a timeout period. It could be caused by either very high network latency or a high packet drop rate.
- **Rec Succ Rate** The receiver success rate for this node since the node statistics were last reset. Receiver success rate is a ratio of the number of packets successfully acknowledged in a first attempt to the total number of received packets. A failure count is incremented when a re-delivery of previously received packet is detected. It could be caused by either very high inbound network latency or lost ACK packets.
- **Member** The member name for this node.
- **Machine** The machine name for this node.
- **Rack** The rack name for this node.
- **Site** The site name for this node.
- **Process** The process name for this node.
- **Uni Addr** The unicast address. This is the IP address of the node's DatagramSocket for point-to-point communication.
- **Uni Port** The unicast port. This is the port of the node's DatagramSocket for point-to-point communication.
- **RoleName** The role name for this node.
- **ProductEdition** The product edition this node is running. Possible values are: Standard Edition (SE), Enterprise Edition (EE), Grid Edition (GE).
- **Send Queue** The number of packets currently scheduled for delivery, including packets sent and still awaiting acknowledgment. Packets that do not receive an acknowledgment within the ResendDelay interval are automatically resent.

#### Packet Transmission Totals

- **Pkts Sent** - Total cumulative packets sent by all nodes in the cluster since the node statistics were last reset.
- **Delta** - Total packets sent by all nodes in the cluster since the last update.
- **Pkts Rcvd** - Total cumulative packets received by all nodes in the cluster since the node statistics were last reset.
- **Delta** - Total packets received by all nodes in the cluster since the last update.
- **Pkts Rptd** - Total cumulative packets repeated by all nodes in the cluster since the node statistics were last reset.
- **Delta** - Total packets repeated by all nodes in the cluster since the last update.
- **Pkts Resent** - Total cumulative packets resent by all nodes in the cluster since the node statistics were last reset.

**Delta** - Total packets resent by all nodes in the cluster since the last update.

### Invocation Service Detail

This display presents detailed information about invocation services. The data displayed here is queried from the Coherence ServiceMBean filtered to only display services of type Invocation. Click on a node in the table to drill down to the "Node Summary" display for that node.

The screenshot shows a window titled "Invocation Service Detail" for cluster "n/a:DemoCluster". It displays a table of invocation services across all hosts. The table has columns for Location, name, Running, CPU %, Messages, Delta, Requests, and Rec. The "JmxNode01.SLHOST1" row is highlighted, showing a CPU usage of 11.5% and 841,145 messages.

Location	name	Running	CPU %	Messages	Delta	Requests	Delta	Rec
ExtendNode01.SLHOST2	Management	✓	0.1	183,053	26	2	0	
ExtendNode01.SLHOST3	Management	✓	0	636,820	30	2	0	
ExtendNode01.SLHOST4	Management	✓	0.3	558,849	48	2	0	
ExtendNode02.SLHOST2	Management	✓	0.5	193,035	50	2	0	
ExtendNode02.SLHOST4	Management	✓	0.8	559,564	50	2	0	
JmxNode01.SLHOST1	Management	✓	11.5	841,145	4,790	839,383	4,790	
OLS-ProcessNode-1.SLHOST2	Management	✓	0.2	207,174	54	2	0	
OLS-ProcessNode-1.SLHOST4	Management	✓	0.2	599,986	52	2	0	
OLS-StoreNode-1.SLHOST2	Management	✓	0.1	263,598	69	2	0	
OLS-StoreNode-1.SLHOST4	Management	✓	0.2	763,897	69	2	0	
ProcessNode01.SLHOST2	Management	✓	0.6	207,681	52	2	0	
ProcessNode01.SLHOST3	Management	✓	0.2	691,792	62	2	0	
ProcessNode01.SLHOST4	Management	✓	0.1	600,603	53	2	0	
ProcessNode02.SLHOST2	Management	✓	0.2	207,370	53	2	0	
ProcessNode02.SLHOST3	Management	✓	0	691,971	59	2	0	
ProcessNode02.SLHOST4	Management	✓	0.2	600,699	53	2	0	
ProcessNode03.SLHOST2	Management	✓	0.3	221,989	57	2	0	
ProcessNode03.SLHOST3	Management	✓	0.3	733,684	66	2	0	
ProcessNode03.SLHOST4	Management	✓	0.5	642,766	59	2	0	
ProcessNode04.SLHOST2	Management	✓	0.3	235,858	61	2	0	
ProcessNode04.SLHOST3	Management	✓	0.5	775,442	67	2	0	
ProcessNode04.SLHOST4	Management	✓	0.2	683,213	62	2	0	
ProcessNode05.SLHOST3	Management	✓	0.2	689,541	59	2	0	
ProcessNode05.SLHOST4	Management	✓	0.3	599,289	52	2	0	
ProcessNode05n.SLHOST2	Management	✓	0.5	235,663	61	2	0	
ProcessNode05n.SLHOST3	Management	✓	0.2	774,342	67	2	0	

**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu, Table open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

**Cluster** Select a cluster to display.

**Host** Select the host for which to display data, or select All Hosts.

#### Invocation Service Information

- **Location** A unique identifier for each node. It is defined as: **member\_name.machine.rack.site**.
- **name** The name of the invocation service.
- **Running** Indicates that the invocation service is running when checked.
- **CPU%** The percent (%) of CPU used by the node.

- **Messages** The number of messages issued by the service to the node in a given time period.
- **Delta** The number of messages received by the node since the last update.
- **Requests** The number of requests issued by the service to the node in a given time period.
- **Delta** The number of requests received by the node since the last update.
- **RequestAverageDuration** The average duration (in milliseconds) of an individual synchronous request issued by the service since the last time the statistics were reset.
- **RequestMaxDuration** The maximum duration (in milliseconds) of a synchronous request issued by the service since the last time the statistics were reset.
- **RequestPendingCount** The number of pending synchronous requests issued by the service.
- **RequestPendingDuration** The duration (in milliseconds) of the oldest pending synchronous request issued by the service.
- **RequestTimeoutCount** The total number of timed-out requests since the last time the statistics were reset.
- **RequestTimeoutMillis** The default timeout value in milliseconds for requests that can be timed-out (e.g. implement the `com.tangosol.net.PriorityTask` interface), but do not explicitly specify the request timeout value.
- **TaskAverageDuration** The average duration (in milliseconds) of an individual task execution.
- **TaskBacklog** The size of the backlog queue that holds tasks scheduled to be executed by one of the service pool threads.
- **TaskCount** The total number of executed tasks since the last time the statistics were reset.
- **TaskHungCount** The total number of currently executing hung tasks.
- **TaskHungDuration** The longest currently executing hung task duration in milliseconds.
- **TaskHungTaskId** The id of the of the longest currently executing hung task.
- **TaskHungThresholdMillis** The amount of time in milliseconds that a task can execute before it is considered hung. Note that a posted task that has not yet started is never considered as hung.
- **TaskMaxBacklog** The maximum size of the backlog queue since the last time the statistics were reset.
- **TaskTimeoutCount** The total number of timed-out tasks since the last time the statistics were reset.
- **TaskTimeoutMillis** The default timeout value in milliseconds for tasks that can be timed-out (e.g. implement the `com.tangosol.net.PriorityTask` interface), but do not explicitly specify the task execution timeout value.
- **ThreadAbandonedCount** The number of abandoned threads from the service thread pool. A thread is abandoned and replaced with a new thread if it executes a task for a period of time longer than execution timeout and all attempts to interrupt it fail.
- **ThreadAverageActiveCount** The average number of active (not idle) threads in the service thread pool since the last time the statistics were reset.
- **ThreadCount** The number of threads in the service thread pool.
- **ThreadIdleCount** The number of currently idle threads in the service thread pool.
- **HostName** Name of the host machine on which the service resides.
- **Throughput** The amount of data (in kilobytes) that is transferred by the service to the node.

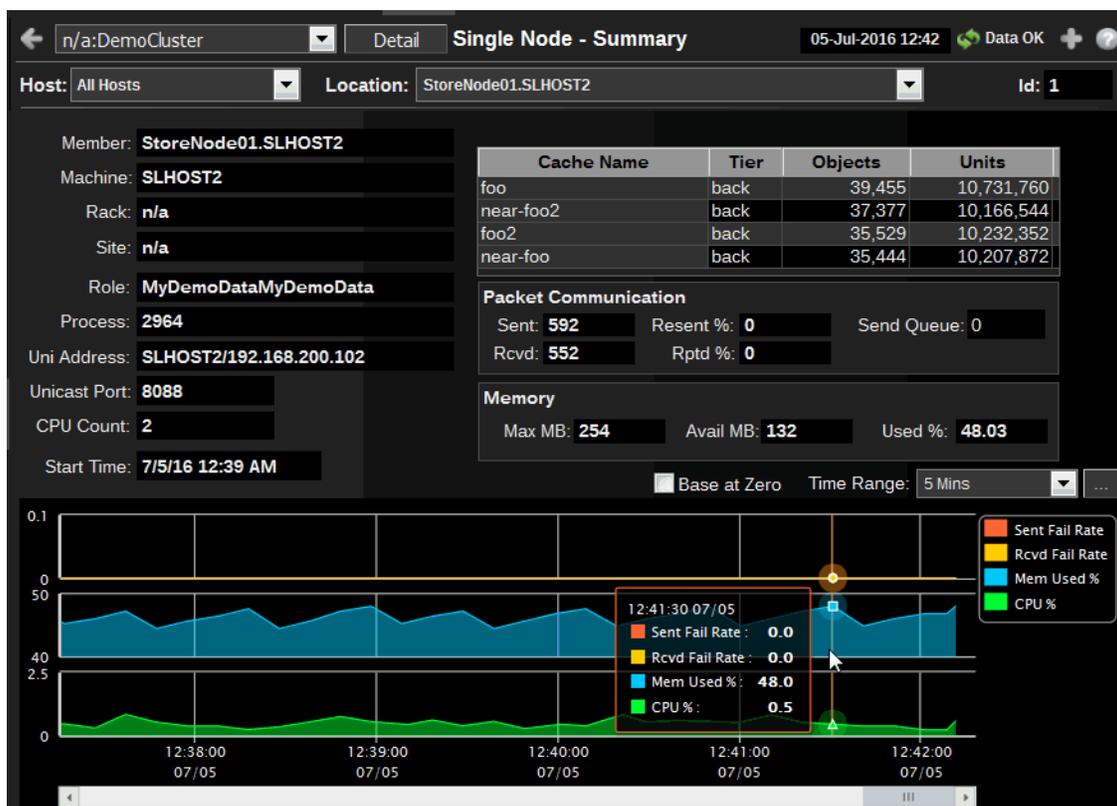
## Single Node

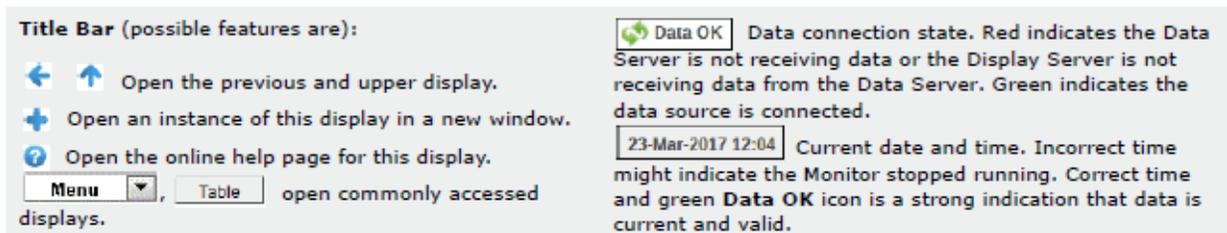
Single Node displays present detailed node performance metrics for a single node. Use the Single Node displays to perform node utilization analysis.

- **"Node Summary"**: Summary view showing details about a single node.
- **"Service Trends"**: Trend graphs showing metrics on a selected node of a selected service. Allows you to visually compare the behavior of metrics over time, for a given node.
- **"Node Detail"**: Tables showing metrics for Node, Cache, Invocation Service, Cache Service, and Storage Manager MBeans.
- **"JVM Summary"**: Runtime, class loader, thread, OS and input arguments.
- **"JVM Memory Trends"**: Heap and non-heap memory trends.
- **"JVM GC Trends"**: Memory usage before and after garbage collection and Garbage Collector activity.
- **"System Properties"**: Table of Java properties for a selected node.

## Node Summary

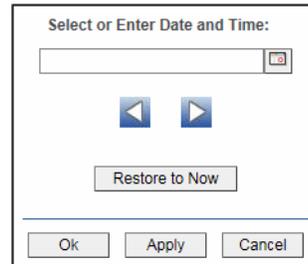
This display presents summary information about an individual node.





<b>Cluster</b>	Select a cluster to display.
<b>Detail</b>	View "Node Detail" display.
<b>Host</b>	Select a host from the drop-down menu.
<b>Location</b>	Select a location from the drop-down menu. <b>Location</b> is a unique identifier for each node and defined as: <b>member_name.machine.rack.site</b> .
<b>Id</b>	The id for the selected node.
<b>Member</b>	The member name for this node.
<b>Machine</b>	The machine name for this node.
<b>Rack</b>	The rack name for this node.
<b>Site</b>	The site name for this node.
<b>Role</b>	The role name for this node.
<b>Process</b>	The process name for this node.
<b>Uni Address</b>	The unicast address. This is the IP address of the node's DatagramSocket for point-to-point communication.
<b>Unicast Port</b>	The unicast port. This is the port of the node's DatagramSocket for point-to-point communication.
<b>CPU Count</b>	Number of CPU cores for the machine this node is running on.
<b>Start Time</b>	The date and time that the selected node joined the cluster.
<b>Cache Data</b>	<p><b>Cache Name</b> Name of Cache.</p> <p><b>Tier</b> Front or Back.</p> <p><b>Objects</b> Number of objects.</p> <p><b>Units</b> Number of units (typically bytes).</p>
<b>Packet Communication</b>	<p><b>Sent</b> Cumulative number of packets sent by this node since the node statistics were last reset.</p> <p><b>Rcvd</b> Cumulative number of packets received by this node since the node statistics were last reset.</p> <p><b>Resent%</b> Cumulative number of packets resent by this node since the node statistics were last reset.</p> <p><b>Rptd%</b> Cumulative number of packets repeated by this node since the node statistics were last reset.</p> <p><b>Send Queue</b> The number of packets currently scheduled for delivery, including packets sent and still awaiting acknowledgment. Packets that do not receive an acknowledgment within the ResendDelay interval are automatically resent.</p>

- Memory**
- Max MB** Total memory allocated.
  - Avail MB** Total memory available.
  - Used%** Percent of allocated memory being used.
- Base at Zero** Use zero as the Y axis minimum for all graph traces.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

- Sent Fail Rate** Percentage of communication packages on this node that failed and needed to be resent.
- Rcvd Fail Rate** Percentage of received communication packages that failed and needed to be repeated.
- Mem Used%** Percent of memory used by the node.
- CPU%** Percent of CPU used by the node.

## Service Trends

Trend graphs showing metrics on a selected node of a selected service. Allows you to visually compare the behavior of metrics over time, for a given node.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

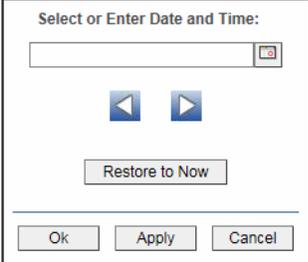
Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

- Cluster** Select a cluster to display.
- Service** Select a service to display.
- Host** Select a host to display.
- Location** Select a location to display. **Location** is a unique identifier for each node and defined as: **member\_name.machine.rack.site**.
- Base at Zero** Use zero as the Y axis minimum for all graph traces.

**Time Range**

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

**Metrics for Service selected by Location**

Trend chart displays the values of labeled Metrics for the selected **Location** over the specified **Time Range**. **Location** is a unique identifier for each node and defined as: **member\_name.machine.rack.site**.

**CPU%** CPU Utilization (as a percent) on the selected **Location** (for example, node).

**Requests** Number of requests issued by the service in the measured period.

**Messages** The number of messages for the given node in the measured interval.

**Request Average** Duration Average duration (in milliseconds) of an individual request issued by the service since the last time the statistics were reset.

**Request Pending** Count Number of pending requests issued by the service.

**Task Backlog** Size of the backlog queue that holds tasks scheduled to be executed by one of the service threads.

**Active Threads** Number of threads in the service thread pool, not currently idle.

## Node Detail

This display presents detailed information about invocation services per node. The data on this display is queried from the Coherence MBeans. NOTE: For details on attributes of these MBeans go to: [http://download.oracle.com/otn\\_hosted\\_doc/coherence/350/com/tangosol/net/management/Registry.html](http://download.oracle.com/otn_hosted_doc/coherence/350/com/tangosol/net/management/Registry.html).

Id	Avail MB	Max MB	Pkts Sent	Delta	Pkts Rcvd	Delta	Pkts Rptd	Delta	Pkts Resent	Delta
1	141	254	3,105,267	729	2,850,223	668	1,094	0	1,256	

Name	Running	CPU %	RequestTotalCount	Requests	Total Messages	Messages	RequestAverageDura
Management	<input checked="" type="checkbox"/>	0.3	2	0	323,611	80	

Service	Running	StatusHA	Storage	CPU %	RequestTotalCount	Requests	Total Messages
DistributedCache	<input checked="" type="checkbox"/>	MACHINE-SAFE	<input checked="" type="checkbox"/>	0.3	378,458	96	1,338,919

Service	Cache Name	Tier	Objects	Hlts	Delta	Misses	Delta
DistributedCache	near-foo	back	35,491	92,130	1	15,671	
DistributedCache	near-foo2	back	37,413	96,278	4	15,329	
DistributedCache	foo2	back	35,587	119,805	12	28,552	

Service	Cache Name	EventsDispatched	EvictionCount	InsertCount	ListenerFilter
DistributedCache	near-foo	0	4,440	39,931	
DistributedCache	foo2	0	6,660	42,247	
DistributedCache	foo	0	38,015	77,698	

### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

- Cluster** Select a cluster to display.
- Summary** View "Node Summary" display.
- Host** Select a host.
- Location** Select a location. **Location** is a unique identifier for each node and defined as: **member\_name.machine.rack.site**.
- Node MBean Data** This table contains data from the Node MBean for the selected node.

**Invocation Service MBean Data**

This table contains data from the Invocation Services MBean for the selected node.

**StatusHA:**

The high availability status of the service:

- ENDANGERED: There is potential data loss in the cluster if a node goes offline.
- NODE-SAFE: There is no risk of data loss in the cluster if a node goes offline (or is taken offline using kill-9). The data is replicated across multiple nodes and remains available in the cluster.
- MACHINE-SAFE: There is no risk of data loss in the cluster if a machine goes offline (or is taken offline using kill-9). The data is replicated across multiple machines and remains available in the cluster.
- RACK-SAFE: There is no risk of data loss in the cluster if a rack goes offline (or is taken offline using kill-9). The data is replicated across multiple racks and remains available in the cluster.
- SITE-SAFE: There is no risk of data loss in the cluster if a site goes offline (or is taken offline using kill-9). The data is replicated across multiple sites and remains available in the cluster.

**Cache Service MBean Data**

This table contains data from the Cache Service and Node MBeans associated with the selected node, as well as the following data.

**Cache MBean Data**

This table contains data from the Cache MBeans associated with the selected node.

**Storage Manager MBean Data**

This table contains data from the Storage Manager MBeans associated with the selected node.

## JVM Summary

Runtime, class loader, thread, OS and input arguments. NOTE: Platform MBean information is available at: [http://java.sun.com/javase/6/docs/api/java/lang/management/package-summary.html#package\\_description](http://java.sun.com/javase/6/docs/api/java/lang/management/package-summary.html#package_description).

The screenshot displays the 'Single Node - JVM Summary' interface. At the top, it shows the cluster 'n/a:DemoCluster' and the selected node 'StoreNode01.SLHOST2' with ID '1'. The interface is divided into several sections:

- Runtime:** Start Time: 7/5/16 12:31 AM; Up Time: 0d 12:15.
- Class Loader:** Loaded Classes: 2333; Unloaded Classes: 77; Total Loaded Classes: 2410.
- Compilation:** Compilation Time: .757 s.
- Threads:** Live Threads: 16; Daemon Threads: 15; Peak Threads: 16.
- Operating System:** Operating System Name: Windows XP; Version: 5.1; Architecture: x86; Available Processors: 2; Percent CPU: 0.5; Total Swap Space Size: 4,096 MB; Free Swap Space Size: 2,630 MB; Total Physical Memory: 2,048 MB; Free Physical Memory: 716 MB; Committed Virtual Memory: 193 MB.

A command line window at the bottom shows the following arguments: `-Xmx256m`, `-Dtangosol.coherence.mbeans=/sl-custom-mbeans.xml`, and `-Dcom.sun.management.jmxremote`.

### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** , **Table** open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

- Cluster** Select a cluster to display.
- Host** Select a host to display.
- Location** Select a location to display. **Location** is a unique identifier for each node and defined as: **member\_name.machine.rack.site**.
- Id** This table contains data from the Node MBean for the selected node.
- Runtime**
  - Start Time** The date and time that the JVM started.
  - Up Time** The uptime of the JVM.

<b>Class Loader</b>	<b>Loaded Classes</b> The number of classes that are currently loaded in the JVM.
	<b>Unloaded Classes</b> The total number of classes unloaded since the JVM started execution.
	<b>Total Loaded Classes</b> The total number of classes that have been loaded since the JVM started execution.
<b>Compilation Time</b>	<p>The approximate accumulated elapsed time (in milliseconds) spent in compilation. If multiple threads are used for compilation, then this value is a summation of the approximate time that each thread spent in compilation.</p> <p>NOTE: Compilation Time monitoring may not be supported depending on the platform (for example, a Java virtual machine implementation).</p>
<b>Threads</b>	<b>Live Threads</b> The number of live threads.
	<b>Daemon Threads</b> The number of live daemon threads.
	<b>Peak Threads</b> The peak live thread count since the Java virtual machine started or peak was reset.
<b>Operating System</b>	<b>Operating System Name</b> The operating system name.
	<b>Version</b> The operating system version.
	<b>Architecture</b> The operating system architecture.
	<b>Available Processors</b> The number of processors available to the JVM.
	<b>Percent CPU</b> Percent of CPU used by the JVM.
	<b>Total Swap Space Size</b> The value of the OperatingSystemMXBean's TotalSwapSpaceSize attribute.
	<b>Free Swap Space Size</b> The value of the OperatingSystem MXBean's FreeSwapSpaceSize attribute.
	<b>Total Physical Memory</b> The value of the OperatingSystemMXBean's TotalPhysicalMemorySize attribute
<b>Free Physical Memory</b> The value of the OperatingSystemMXBean's FreePhysicalMemorySize attribute	
<b>Committed Virtual Memory</b> The value of the OperatingSystemMXBean's CommittedVirtualMemorySize attribute	
<b>Input Arguments</b>	The list of JVM arguments in the RuntimeMXBean's InputArguments attribute.

## JVM Memory Trends

Heap and non-heap memory trends. NOTE: Platform MBean information is available at: [http://java.sun.com/javase/6/docs/api/java/lang/management/package-summary.html#package\\_description](http://java.sun.com/javase/6/docs/api/java/lang/management/package-summary.html#package_description).



### Title Bar (possible features are):

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu, Table open commonly accessed displays.

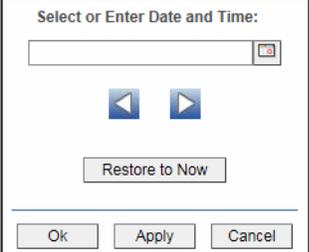
Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

- Cluster** Select a cluster to display.
- Host** Select a host to display.
- Location** Select a location to display. **Location** is a unique identifier for each node and defined as: **member\_name.machine.rack.site**.
- Id** This table contains data from the Node MBean for the selected node.
- Base at Zero** Use zero as the Y axis minimum for all graph traces.

**Time Range**

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

**Heap Memory**

**Maximum** The value of the max field within the MemoryMXBean HeapMemoryUsage attribute.

**Committed** The value of the committed field within the MemoryMXBean HeapMemoryUsage attribute.

**Used** The value of the used field within the MemoryMXBean HeapMemoryUsage attribute.

**Peak Tenured Used** The value of the used field within the TenuredGen MemoryPoolMXBean PeakUsage attribute.

**Non-Heap Memory**

**Maximum** The value of the max field within the MemoryMXBean NonHeapMemoryUsage attribute.

**Committed** The value of the committed field within the MemoryMXBean NonHeapMemoryUsage attribute.

**Used** The value of the used field within the MemoryMXBean NonHeapMemoryUsage attribute.

**Objects Pending Finalization** The value of the MemoryMXBean ObjectPendingFinalizationCount attribute.

**Verbose** The value of the MemoryMXBean Verbose attribute.

**Garbage Collection**

**name** Name of the Garbage Collector MBean.

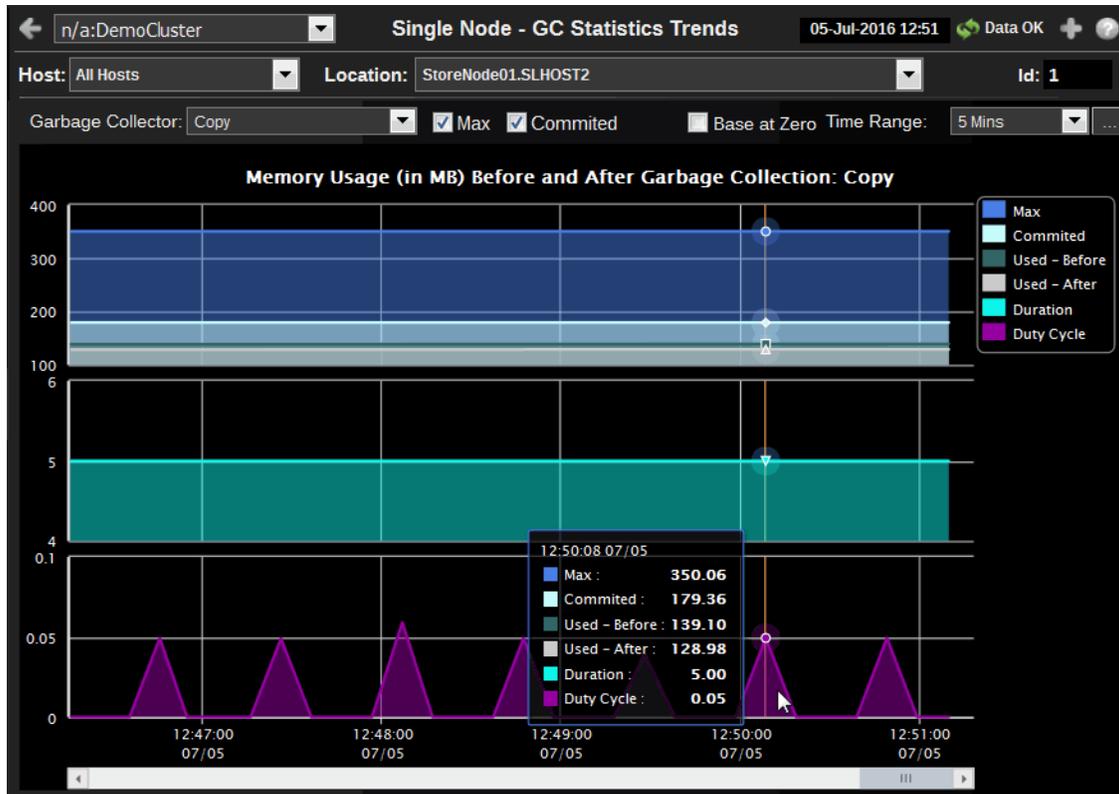
**LastGcInfo.Count** The GcThreadCount from the Garbage Collector's LastGcInfo MBean.

**LastGcInfo.Duration** The Duration from the Garbage Collector's LastGcInfo MBean.

**Operations Run Garbage Collector** Executes the MemoryMXBean garbage collection operation, **Reset Peak Usage** Executes the TenuredGen resetPeakUsage operation.

## JVM GC Trends

Memory usage before and after garbage collection and Garbage Collector activity. NOTE: Platform MBean information is available at: [http://java.sun.com/javase/6/docs/api/java/lang/management/package-summary.html#package\\_description](http://java.sun.com/javase/6/docs/api/java/lang/management/package-summary.html#package_description).



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

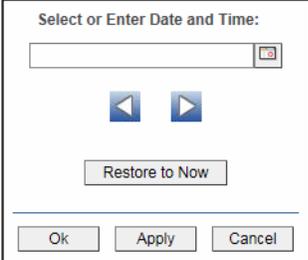
- Cluster** Select a cluster to display.
- Host** Select a host to display.
- Location** Select a location to display. **Location** is a unique identifier for each node and defined as: **member\_name.machine.rack.site**.
- Id** This table contains data from the Node MBean for the selected node.
- Garbage Collector** Select a Garbage Collector.
- Max** Select to add the Max trace (graph will rescale if necessary).
- Committed** Select to add the Committed trace (graph will rescale if necessary).

**Base at Zero**

Use zero as the Y axis minimum for all graph traces.

**Time Range**

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

**Memory Usage (in MB) Before and After Garbage Collection**

**Max** The maximum amount of memory used by the node or nodes.

**Committed** The amount of memory guaranteed to be available for use by the JVM.

**Used - Before** The amount of memory used by the node or nodes before garbage collection.

**Used - After** The amount of memory used by the node or nodes after garbage collection.

**Duration** The duration, in seconds, that memory is used by the node or nodes.

**Duty** Cycle Percent of time spent by the node or nodes in garbage collection.

## System Properties

Table of Java properties for a selected node.

key	value	Connection
awt.toolkit	sun.awt.windows.WToolkit	DemoCluste
com.sun.management.jmxremote	true	DemoCluste
file.encoding	Cp1252	DemoCluste
file.encoding.pkg	sun.io	DemoCluste
file.separator	\	DemoCluste
java.awt.graphicsenv	sun.awt.Win32GraphicsEnvironment	DemoCluste
java.awt.printerjob	sun.awt.windows.WPrinterJob	DemoCluste
java.class.path	.;C:\rtvdemos\rtvoc_57c1\conf;C:\rtvdemos\rtvoc_57c1	DemoCluste
java.class.version	50.0	DemoCluste
java.endorsed.dirs	C:\Program Files\Java\jre6\lib\endorsed	DemoCluste
java.ext.dirs	C:\Program Files\Java\jre6\lib\ext;C:\WINDOWS\Sun\Java	DemoCluste
java.home	C:\Program Files\Java\jre6	DemoCluste
java.io.tmpdir	C:\DOCUME~1\LOCALS~1\Temp\	DemoCluste
java.library.path	C:\WINDOWS\system32;.;C:\WINDOWS\Sun\Java\bin;C	DemoCluste
java.rmi.server.randomIDs	true	DemoCluste
java.runtime.name	Java(TM) SE Runtime Environment	DemoCluste
java.runtime.version	1.6.0_11-b03	DemoCluste
java.specification.name	Java Platform API Specification	DemoCluste
java.specification.vendor	Sun Microsystems Inc.	DemoCluste
java.specification.version	1.6	DemoCluste
java.vendor	Sun Microsystems Inc.	DemoCluste
java.vendor.url	http://java.sun.com/	DemoCluste
java.vendor.url.bug	http://java.sun.com/cgi-bin/bugreport.cgi	DemoCluste
java.version	1.6.0_11	DemoCluste

### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

<b>Cluster</b>	Select a cluster to display.
<b>Host</b>	Select a host to display.
<b>Location</b>	Select a location to display. <b>Location</b> is a unique identifier for each node and defined as: <b>member_name.machine.rack.site</b> .
<b>Id</b>	This table contains data from the Node MBean for the selected node.
<b>java.runtime.version</b>	The value of the RuntimeMXBeans's VmVersion attribute.
<b>System Properties</b>	This table contains the attribute/value pairs from the RuntimeMXBean's SystemProperties attribute.

## Time Range Analysis

These displays allow you to compare data between two sets of time ranges.

- “Service Comparison” on page 799: Analyze service data for two sets of time ranges.
- “Cache Comparison” on page 800: Analyze cache data for two sets of time ranges.

### Service Comparison

This display allows for analysis of service data for two sets of time ranges.

Service: DistributedCache

Time Range Analysis:

Time Range 1: Start: [ ] End: [ ]

Time Range 2: Start: [ ] End: [ ]

Metric Name	Time Range 1 Value	Time Range 2 Value	Percentage Change
DeltaMessages	444,734,253	444,734,253	0.00
DeltaRequestTotalCount	123,883,421	123,883,421	0.00
DeltaTaskCount	0	0	0.00
TaskBacklog	22,701	22,701	0.00

#### Title Bar (possible features are):

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu Table open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

**Cluster** Select a cluster to display.

**Service** Select a service to display.

**Storage Nodes** Select to display storage node data in the trend graphs of this display.

**Process Nodes** Select to display process node data in the trend graphs of this display.

**Time Range Analysis**

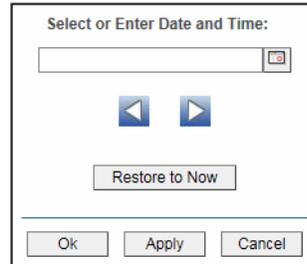
**Time Range 1:** Set Start and End times for Time Range 1

**Time Range 2:** Set Start and End times for Time Range 2

Time Range 1: Data Bucket Timestamp and Time Range 2: Data Bucket Timestamp displays the Start and End timestamps for the actual data buckets used in the comparison, since data may be compacted into buckets with different Start and End times from the specified values.

**Time Range**

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. **NOTE:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

**Cache Comparison**

This display allows for analysis of cache data for two sets of time ranges.

Metric Name	Time Range 1 Value	Time Range 2 Value	Percentage Change
DeltaTotalGets	88,477,277	88,477,277	0.00
DeltaTotalPuts	51,678,742	51,678,742	0.00
DeltaCacheHits	62,646,428	62,646,428	0.00
DeltaStoreReads	0	0	
DeltaStoreWrites	0	0	
DeltaStoreReadMillis	0	0	
DeltaStoreWriteMillis	0	0	
DeltaStoreFailures	0	0	

**Title Bar** (possible features are):

-   Open the previous and upper display.
-  Open an instance of this display in a new window.
-  Open the online help page for this display.

open commonly accessed displays.

 **Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Cluster** Select a cluster to display.

**Service** Select a service to display.

**Storage Nodes** Select to display storage node data in the trend graphs of this display.

**Process Nodes** Select to display process node data in the trend graphs of this display.

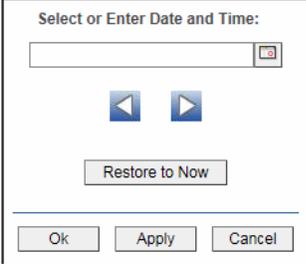
**Time Range Analysis**

**Time Range 1:** Set Start and End times for Time Range 1

**Time Range 2:** Set Start and End times for Time Range 2

Time Range 1: Data Bucket Timestamp and Time Range 2: Data Bucket Timestamp displays the Start and End timestamps for the actual data buckets used in the comparison, since data may be compacted into buckets with different Start and End times from the specified values.

**Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



The dialog box titled "Select or Enter Date and Time:" contains a text input field at the top with a calendar icon on the right. Below the input field are two navigation arrows: a left-pointing arrow and a right-pointing arrow. Underneath the arrows is a button labeled "Restore to Now". At the bottom of the dialog are three buttons: "Ok", "Apply", and "Cancel".

By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

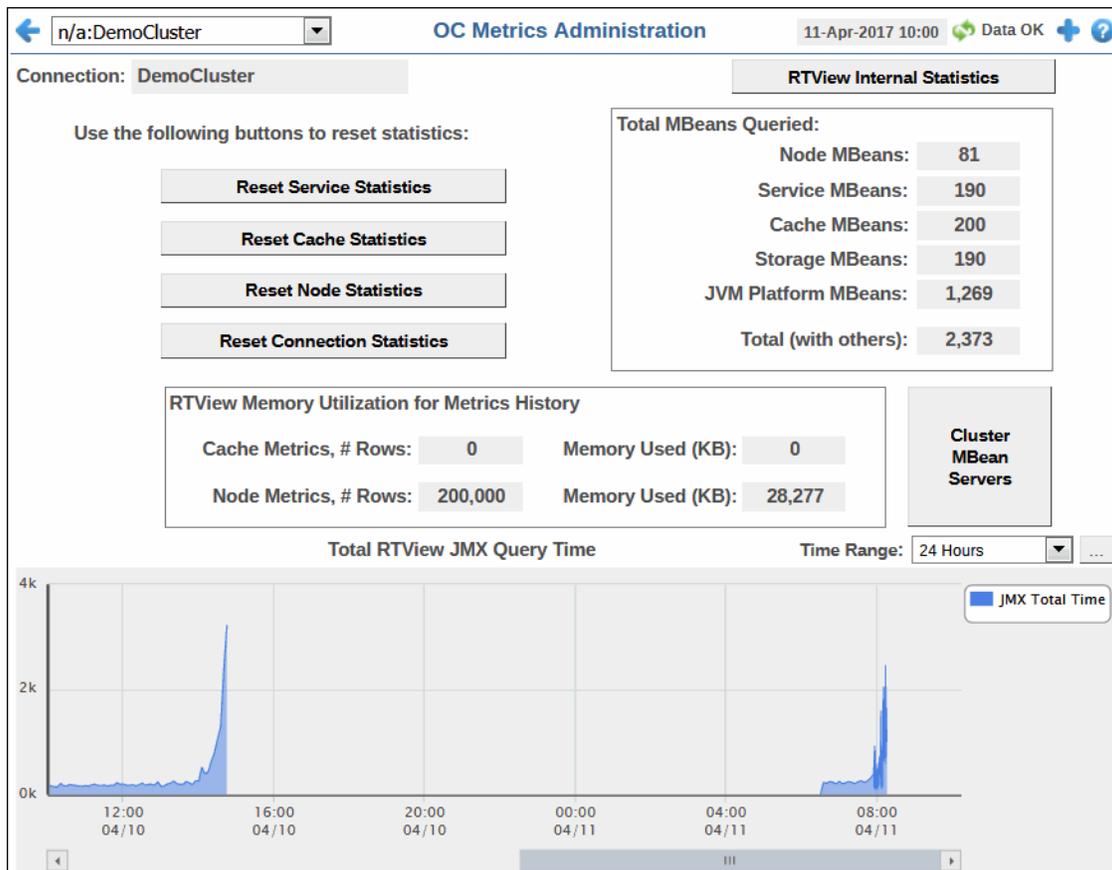
## OC Administration

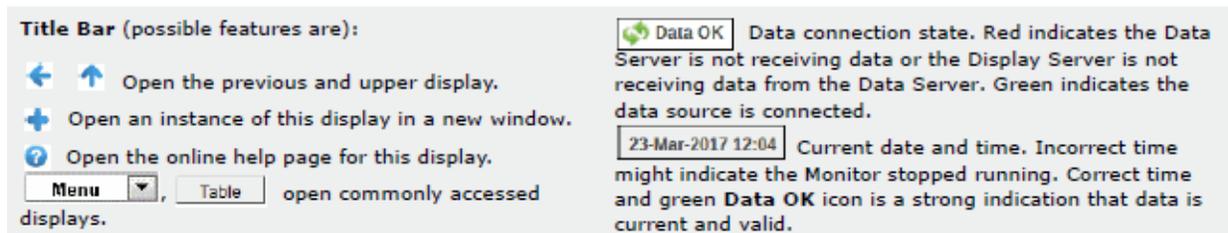
These displays allow you to manage your Oracle Coherence metrics, nodes and caches. Some of these displays might be read-only depending on your login credentials.

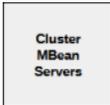
- ["OC Metrics Administration" on page 802](#): Monitor information on metrics acquisition. Permits user to reset system metrics.
- ["Cluster MBean Servers" on page 804](#): Access this display using the **Cluster MBean Servers** button in the ["OC Metrics Administration"](#) display. Permits user to find and choose a different MBean server.
- ["Management Settings" on page 806](#): Monitor information about Coherence JMX management settings.
- ["Node Administration" on page 808](#): Permits user to modify node parameters.
- ["Cache Administration" on page 810](#): Permits user to modify cache parameters.

## OC Metrics Administration

This display allows various statistics to be reset, so that cumulative data can be visualized more meaningfully. It is read-only unless you are logged in as admin or super.

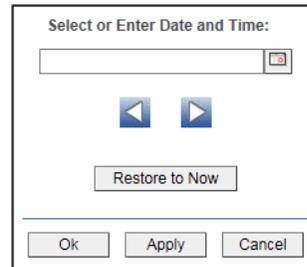




<b>Cluster</b>	Select a cluster to display.
<b>Connection</b>	The name of the JMX connection used to access the cluster data.
<b>Reset Service Statistics</b>	Click to reset the cumulative counts of the service statistics.
<b>Reset Cache Statistics</b>	Click to reset the cumulative counts of the cache statistics.
<b>Reset Node Statistics</b>	Click to reset the cumulative counts of the node statistics.
<b>Reset Connection Statistics</b>	Click to reset the cumulative counts of the connection statistics.
<b>Total MBeans Queried</b>	<p><b>Node MBeans</b> Total number of node MBeans queried.</p> <p><b>Service MBeans</b> Total number of service MBeans queried.</p> <p><b>Cache MBeans</b> Total number of cache MBeans queried.</p> <p><b>Storage MBeans</b> Total number of storage MBeans queried.</p> <p><b>JVM Platform MBeans</b> Total number of JVM platform MBeans queried.</p> <p><b>Total</b> Total number of MBeans queried.</p>
<b>RTView Memory Utilization for Metrics History</b>	<p>By default, the Oracle Coherence Monitor stores several hours of data using in-memory tables.</p> <p><b>Cache Metrics, # Rows</b> The number of table rows used by the Monitor to store cache metrics data.</p> <p><b>Cache Metrics, Memory Used (KB)</b> The amount of memory (KB) used by the Monitor to store cache metrics data.</p> <p><b>Node Metrics, # Rows</b> The number of table rows used by the Monitor to store node metrics data.</p> <p><b>Node Metrics, Memory Used (KB)</b> The amount of memory (KB) used by the Monitor to store node metrics data.</p>
<b>Cluster MBean Servers</b>	Click to open the " <a href="#">Cluster MBean Servers</a> " display which lists the currently detected remote JMX management enabled MBean Servers in the selected cluster. If your MBean server goes down, use this display to find and choose a different available MBean server.
	

**Total RTView JMX Query Time** Traces the total amount of time, in milliseconds, to query the monitoring MBeans from Coherence.

**Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

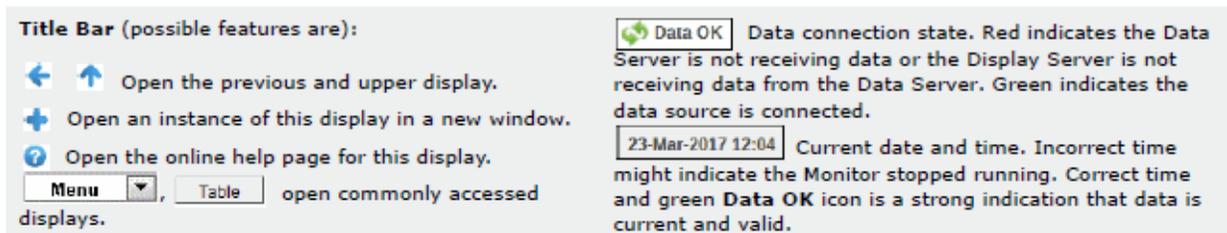
Click **Restore to Now** to reset the time range end point to the current time.

## Cluster MBean Servers

Access this display using the **Cluster MBean Servers** button in the "OC Metrics Administration" display.

View a list of URLs for all currently detected remote JMX management enabled MBean Servers in a cluster. Information displayed includes the hostname and IP address of the cluster node, and the port used for remote JMX management.

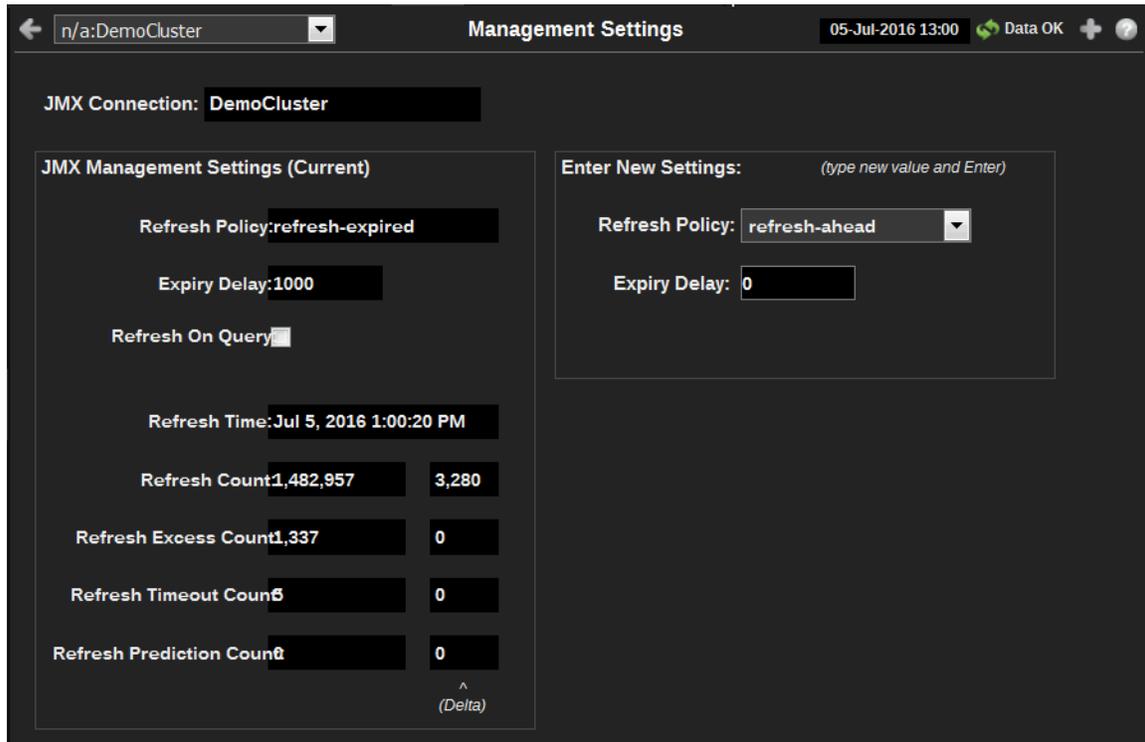
Cluster MBean Servers						
Connection	nodeId	HostName	IpAddress	port	Location	InputArguments
DemoCluster	30	localhost	0.0.0.0	9971	JmxNode01.SLHOST1	-Xmx256m;-Dswin



<b>Cluster</b>	Select a cluster to display.
<b>Connection</b>	The name of the JMX connection used to access the cluster data.
<b>nodeId</b>	The unique identifier for the MBean Server.
<b>HostName</b>	The name of the host for the MBean Server.
<b>IpAddress</b>	The IP address for the MBean Server.
<b>port</b>	The port number for the MBean Server.
<b>Location</b>	A unique identifier for each node. It is defined as: <b>member_name.machine.rack.site</b> .
<b>InputArguments</b>	A list of JVM arguments in the Runtime JMX MBean's InputArguments attribute.
<b>Expired</b>	When checked, this connection is expired due to inactivity.

## Management Settings

This display is read-only unless you are logged in as admin or super.



**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

**Cluster** Select a cluster to display.

**JMX Connection** The name of the JMX connection used to access the cluster data.

**JMX Management Settings****Refresh Policy:**

Select a refresh policy from the drop-down list.

**refresh-expired** Each MBean will be refreshed from the remote node when it is accessed and the expiry delay has passed from the last refresh (same functionality as in pre-3.4 Coherence releases. This option is the default setting and is best used when MBeans are accessed in a random pattern.

**refresh-ahead** MBeans are refreshed before they are requested based on prior usage patterns after the expiry delay has passed, reducing latency of management information with a minor increase in network consumption. This option is best when MBeans are accessed in a repetitive/programmatic pattern.

**refresh-behind** Each MBean will be refreshed after the data is accessed, ensuring optimal response time. However, note that the information returned will be offset by the last refresh time.

**refresh-onquery** Select this option if the refresh-on-query MBeanServer is configured.

**Expiry Delay:**

Duration (in milliseconds) that the MBeanServer will keep a remote model snapshot before refreshing.

**Refresh on Query:**

Specifies whether or not the refresh-on-query MBeanServer is configured. If so, then set the RefreshPolicy to refresh-onquery.

**Refresh Time** The timestamp when this model was last retrieved from a corresponding node. For local servers it is the local time.

**Refresh Count\*** The total number of snapshots retrieved since the statistics were last reset.

**Refresh Excess Count** The number of times the MBean server predictively refreshed information and the information was not accessed. Delta values show the change in the counts within the most recent JMX retrieval period.

**Refresh Timeout Count\*** The number of times this management node has timed out while attempting to refresh remote MBean attributes.

**Refresh Prediction Count\*** The number of times the MBeanServer used a predictive (refresh-behind, refresh-ahead, refresh-onquery) algorithm to refresh MBean information.

## Node Administration

This display allows the user to view and change settings for individual Nodes. It is read-only unless you are logged in as super. Click on the desired Node to select that Node. Change the data item in the bottom half of the display and press Return to make the change. All data on this display is queried from and set on the Coherence ClusterNodeMBean.

The screenshot shows the 'Node Administration' window for a cluster named 'n/a:DemoCluster'. The window title is 'Node Administration' and it shows the date and time '05-Jul-2016 13:01' along with a 'Data OK' status icon. Below the title bar, there is a section titled 'Current Settings for all Nodes (Select a Node to change settings)'. This section contains a table with the following columns: Location, BufferPublishSize, BufferReceiveSize, BurstCount, BurstDelay, and MulticastThreshold. The table lists various nodes such as 'ExtendNode01.SLHOST4', 'ExtendNode02.SLHOST2', 'JmxNode01.SLHOST1', and 'ProcessNode01.SLHOST3'. Below the table, there are controls for 'Selected Node(s): All' and a 'Select All Nodes' button. The bottom half of the window contains configuration options for the selected node, including 'Burst Count: 0', 'Burst Delay: 0', 'Logging Level: 0', 'Logging Limit: 0', 'Logging Format: [text box]', 'Shutdown Node', 'Ensure Cache Service: DistributedCache', 'Ensure Invocation Service: [text box]', 'Multicast Threshold: 0', 'Resend Delay: 0', 'Send Ack Delay: 0', 'Traffic Jam Count: 0', and 'Traffic Jam Delay: 0'.

Location	BufferPublishSize	BufferReceiveSize	BurstCount	BurstDelay	MulticastThreshold
ExtendNode01.SLHOST4	32	1428	0	10	25
ExtendNode02.SLHOST2	32	1428	0	10	25
ExtendNode02.SLHOST4	32	1428	0	10	25
JmxNode01.SLHOST1	32	1428	0	10	25
OLS-ProcessNode-1.SLHOST2	32	1428	0	10	25
OLS-ProcessNode-1.SLHOST4	32	1428	0	10	25
OLS-StoreNode-1.SLHOST2	32	1428	0	10	25
OLS-StoreNode-1.SLHOST4	32	1428	0	10	25
ProcessNode01.SLHOST2	32	1428	0	10	25
ProcessNode01.SLHOST3	32	1428	0	10	25
ProcessNode01.SLHOST4	32	1428	0	10	25

### Title Bar (possible features are):

- ← ↑ Open the previous and upper display.
- ⊕ Open an instance of this display in a new window.
- 🔍 Open the online help page for this display.
- Menu ▾, Table open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

**Cluster** Select a cluster to display.

**JMX Connection** The name of the JMX connection used to access the cluster data.

**Current Settings  
for All Nodes**

**Location** A unique identifier for each node. It is defined as:  
**member\_name.machine.rack.site.**

**BufferPublishSize** The buffer size of the unicast datagram socket used by the Publisher, measured in the number of packets. Changing this value at runtime is an inherently unsafe operation that will pause all network communications and may result in the termination of all cluster services.

**BufferReceiveSize** The buffer size of the unicast datagram socket used by the Receiver, measured in the number of packets. Changing this value at runtime is an inherently unsafe operation that will pause all network communications and may result in the termination of all cluster services.

**BurstCount** The maximum number of packets to send without pausing. Anything less than one (e.g. zero) means no limit.

**BurstDelay** The number of milliseconds to pause between bursts. Anything less than one (e.g. zero) is treated as one millisecond.

**MulticastThreshold** The percentage (0 to 100) of the servers in the cluster that a packet will be sent to, above which the packet will be multicasted and below which it will be unicasted.

**ResendDelay** The minimum number of milliseconds that a packet will remain queued in the Publisher's re-send queue before it is resent to the recipient(s) if the packet has not been acknowledged. Setting this value too low can overflow the network with unnecessary repetitions. Setting the value too high can increase the overall latency by delaying the re-sends of dropped packets. Additionally, change of this value may need to be accompanied by a change in SendAckDelay value.

**SendAckDelay** The minimum number of milliseconds between the queueing of an Ack packet and the sending of the same. This value should be not more than a half of the ResendDelay value

**TrafficJamCount** The maximum total number of packets in the send and resend queues that forces the publisher to pause client threads. Zero means no limit.

**TrafficJamDelay** The number of milliseconds to pause client threads when a traffic jam condition has been reached. Anything less than one (e.g. zero) is treated as one millisecond.

**LoggingLevel** Specifies which logged messages will be output to the log destination. Valid values are non-negative integers or -1 to disable all logger output.

**LoggingLimit** The maximum number of characters that the logger daemon will process from the message queue before discarding all remaining messages in the queue. Valid values are integers in the range [0...]. Zero implies no limit.

**LoggingFormat** Specifies how messages will be formatted before being passed to the log destination

**LoggingDestination** The output device used by the logging system. Valid values are **stdout**, **stderr**, **jdk**, **log4j**, or a file name.

**nodeId** The short Member id that uniquely identifies the Member at this point in time and does not change for the life of this Member.

**ProcessName** A configured name that should be the same for Members that are in the same process (JVM), and different for Members that are in different processes. If not explicitly provided, for processes running with JRE 1.5 or higher the name will be calculated internally as the Name attribute of the system RuntimeMXBean, which normally represents the process identifier (PID).

**Selected Node(s)** Lists the nodes selected in the table.

**Select All Nodes** Click to select all nodes.

**Shutdown Node** Stop all the clustered services running at this node (controlled shutdown). The management of this node will not be available until the node is restarted (manually or by programming).

**Ensure Cache Service**

Ensure that a CacheService for the specified cache runs at the cluster node represented by this MBean. This method will use the configurable cache factory to find out which cache service to start if necessary. Return value indicates the service name; null if a match could not be found.

**Ensure Invocation**

Ensure that an InvocationService with the specified name runs at the cluster node represented by this MBean.

**Cache Administration**

This display allows the user to view and change settings for individual caches. It is read-only unless you are logged in as super. Click on the desired cache to select that cache. Change the data item in the bottom half of the display and press Return to make the change. The data on this display is queried from and set on the Coherence CacheMBean.

The screenshot shows the 'Cache Administration' window for a cluster named 'n/a:DemoCluster'. The 'Service' is set to 'DistributedCache' and the 'Cache' is 'foo'. A table lists various nodes with columns for Location, LowUnits, HighUnits, BatchFactor, ExpiryDelay, FlushDelay, and Queue. Below the table, there are input fields for 'High Units', 'Low Units', 'Batch Factor', 'Expiry Delay', 'Flush Delay', 'Queue Delay', 'Refresh Factor', and 'Requeue Threshold', all currently set to 0. A 'Data OK' icon is visible in the top right corner.

Location	LowUnits	HighUnits	BatchFactor	ExpiryDelay	FlushDelay	Queue
StoreNode01.SLHOST1	8,000,000	12,000,000	0.0	0	0	0
StoreNode01.SLHOST2	8,000,000	12,000,000	0.0	0	0	0
StoreNode01.SLHOST4	8,000,000	12,000,000	0.0	0	0	0
StoreNode04.SLHOST1	8,000,000	12,000,000	0.0	0	0	0
StoreNode04.SLHOST4	8,000,000	12,000,000	0.0	0	0	0
StoreNode05.SLHOST2	8,000,000	12,000,000	0.0	0	0	0
StoreNode05.SLHOST4	8,000,000	12,000,000	0.0	0	0	0
StoreNode05n.SLHOST1	8,000,000	12,000,000	0.0	0	0	0
StoreNode05n.SLHOST2	8,000,000	12,000,000	0.0	0	0	0
StoreNode05n.SLHOST4	8,000,000	12,000,000	0.0	0	0	0

**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

- Cluster** Select a cluster to display.
- Service** Select the service to display.
- Cache** Select the cache to display.

<b>Front</b>	Select for front tier, deselect for back tier.
<b>Current Settings for all Nodes on Selected Cache</b>	<p><b>Location</b> A unique identifier for each node. It is defined as: <b>member_name.machine.rack.site.</b></p> <p><b>LowUnits</b> The number of units to which the cache will shrink when it prunes. This is often referred to as a `low water mark` of the cache.</p> <p><b>HighUnits</b> The limit of the cache size measured in units. The cache will prune itself automatically once it reaches its maximum unit level. This is often referred to as the `high water mark` of the cache.</p> <p><b>BatchFactor</b> The BatchFactor attribute is used to calculate the `soft-ripe` time for write-behind queue entries. A queue entry is considered to be `ripe` for a write operation if it has been in the write-behind queue for no less than the QueueDelay interval. The `soft-ripe` time is the point in time prior to the actual `ripe` time after which an entry will be included in a batched asynchronous write operation to the CacheStore (along with all other `ripe` and `soft-ripe` entries). This attribute is only applicable if asynchronous writes are enabled (for example, the value of the QueueDelay attribute is greater than zero) and the CacheStore implements the storeAll() method. The value of the element is expressed as a percentage of the QueueDelay interval. Valid values are doubles in the interval [0.0, 1.0].</p> <p><b>ExpiryFactor</b> The time-to-live for cache entries in milliseconds. Value of zero indicates that the automatic expiry is disabled. Change of this attribute will not affect already-scheduled expiry of existing entries.</p> <p><b>FlushDelay</b> The number of milliseconds between cache flushes. Value of zero indicates that the cache will never flush.</p> <p><b>QueueDelay</b> The number of seconds that an entry added to a write-behind queue will sit in the queue before being stored via a CacheStore. Applicable only for WRITE-BEHIND persistence type.</p> <p><b>RefreshFactor</b> The RefreshFactor attribute is used to calculate the `soft-expiration` time for cache entries. Soft-expiration is the point in time prior to the actual expiration after which any access request for an entry will schedule an asynchronous load request for the entry. This attribute is only applicable for a ReadWriteBackingMap which has an internal LocalCache with scheduled automatic expiration. The value of this element is expressed as a percentage of the internal LocalCache expiration interval. Valid values are doubles in the interval[0.0, 1.0]. If zero, refresh-ahead scheduling will be disabled.</p> <p><b>Requeue Threshold</b> The maximum size of the write-behind queue for which failed CacheStore write operations are requeued. If zero, the write-behind requeueing will be disabled. Applicable only for WRITE-BEHIND persistence type.</p> <p><b>nodeId</b> The node ID.</p>
<b>Selected Node(s)</b>	Lists the nodes selected in the table.
<b>Select All Nodes</b>	Click to select all nodes in the table.

## RTView Cache Tables

View data that RTView is capturing and maintaining. Drill down and view details of RTView Cache Tables. Use this data for debugging. This display is typically used for troubleshooting with Technical Support.

Choose a cache table from the upper table to see cached data.

**RTView Cache Tables** 05-Jul-2016 13:09 Data OK History Tables

DataServer: <Default> Max Rows: 4000

CacheTable	TableType	Rows	Columns	Memory
JmxStatsTotals	current	1	4	441
OcBadCommunicationNodes	current	140	6	14,999
OcCacheServiceStats	current	88	58	62,666
OcCacheServiceTotals	current	8	26	4,441
OcCacheStats	current	172	80	206,148
OcCacheTotals	current	17	52	13,406
OcClusterOverview	current	1	7	791
OcClusterStats	current	1	19	14,103
OcExtendConnections	current	112	30	68,304
<b>OclnvoationServiceStats</b>	current	<b>63</b>	<b>60</b>	<b>62,252</b>
OclnvoationServiceTotals	current	1	26	2,841
OcJmxConnection	current	2	7	1,254
OcJmxHostData	current	1	15	1,754

**OclnvoationServiceStats** Rows: 63

TIME_STAMP	BackupCou	OwnedParti	OwnedParti	RefreshTim	RequestAvc	RequestMa	RequestPei	RequestPei	Request
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2016 1	85.5	109	0	0	
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2016 1	8	16	0	0	
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2016 1	8	16	0	0	
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2016 1	8	16	0	0	
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2016 1	7.5	15	0	0	
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2016 1	16	16	0	0	
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2016 1	15.5	16	0	0	
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2016 1	24	32	0	0	
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2016 1	31.5	47	0	0	
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2016 1	39.5	63	0	0	
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2016 1	8	16	0	0	
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2016 1	16	16	0	0	
07/05/16 13:09:06	-1	-1	-1	Jul 5, 2016 1	8	16	0	0	

**Title Bar** (possible features are):

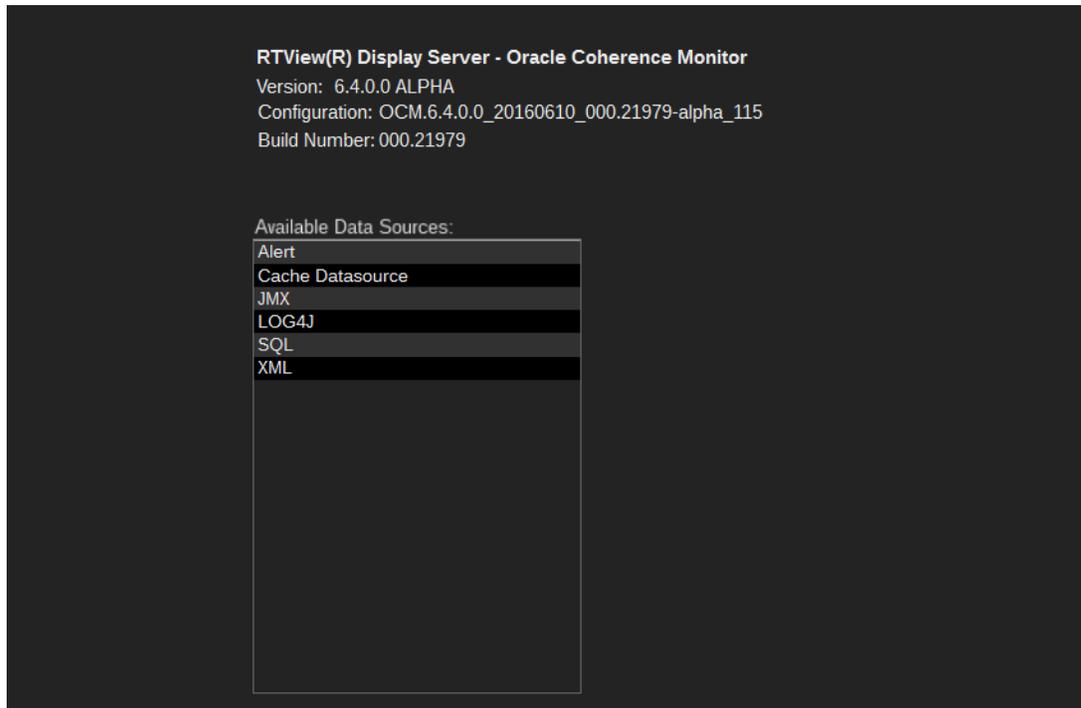
- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu , Table open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

## About

This display shows details about the Solution Package version and data sources available to your system.



---

## Oracle Coherence - HTML

The Oracle Database HTML displays provide extensive visibility into the health and performance of Oracle databases. The HTML version features an ["Oracle Database Overview Display- HTML"](#) (pictured below), and several Views which can be found under **Components** tab > **Databases** > **Oracle Database**:

- ["Oracle Database Overview Display- HTML"](#): This display presents a health snap-shot of the Oracle Database system.
- ["Oracle Databases View - HTML"](#): These displays present metrics about Oracle databases.
- ["Oracle Instances View - HTML"](#): These displays present metrics about instances on Oracle databases.Si

## Oracle Database Overview Display- HTML

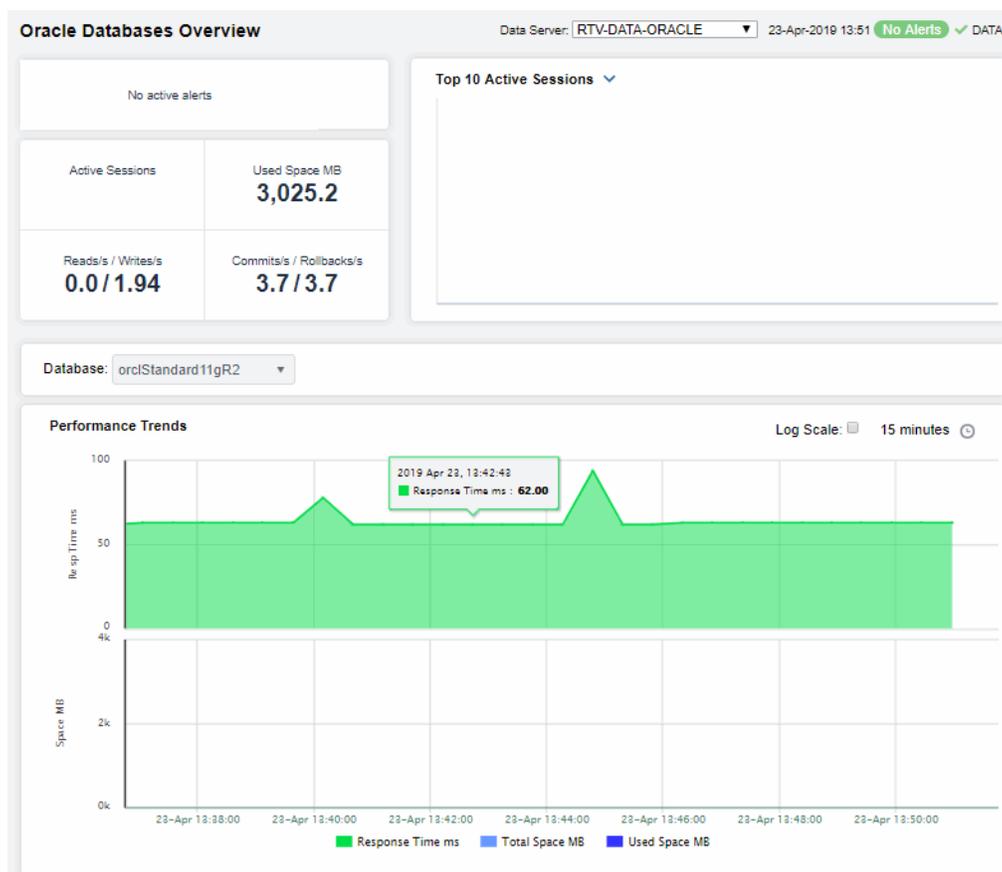
The **Oracle Database Overview** is the top-level display for the Oracle Database Solution Package, which provides a good starting point for immediately getting the status of all your Oracle Databases on your Data Server.

You can select the RTView DataServer for which you want to see data and easily view the current data for that DataServer including:

- The total number of **Active Alerts**, including the total number of critical and warning alerts.
- The number of **Active Sessions** across all databases.
- The total amount of **Used Space** across all databases.
- The number of **Reads / Writes** per second across all databases.
- The number of **Commits / Rollbacks** per second across all databases.
- A bar graphs showing the **Top 10 Active Sessions**.

You can hover over each region in the upper half of the Overview to see more detail. You can also drill-down to see even more detail in the "[Oracle Databases Table - HTML](#)" by clicking on each respective region in the Overview.

The bottom half of the display allows you to select a database for the performance trend graph to trace **Response Time**, **Total Space** and **Used Space**. You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.



## Oracle Databases View - HTML

These displays present performance metrics for Oracle databases. Displays are:

- ["Oracle Databases Table - HTML"](#): A list of all Oracle databases with detailed utilization metrics.
- ["Databases Heatmap - HTML"](#): A heatmap shows alert status of all Oracle databases.
- ["Databases Summary - HTML"](#): Detailed performance metrics for one Oracle database.

### Oracle Databases Table - HTML

Investigate and compare detailed utilization metrics for all Oracle databases. This display contains all metrics available for Oracle databases, including **Space Used**, **Free Space**, **Total Space** and the number of **Instances** on each.

Each row in the table contains data for a particular Oracle database. Click a column header to sort column data in ascending or descending order. Double-click on a table row to drill-down to the ["Databases Summary - HTML"](#) display and view metrics for that particular Oracle database. Toggle between the commonly accessed Heatmap and Summary displays by clicking the drop down list on the display title.

Right-click on a table cell to **Export to Excel** or **Copy Cell Value**.

The **Alert Level** column indicates the most critical alert on the database, where:

- Red indicates that one or more alerts exceeded their ALARM LEVEL threshold in the table row.
- Yellow indicates that one or more alerts exceeded their WARNING LEVEL threshold in the table row.
- Green indicates that no alerts have reached their alert thresholds.

Oracle Databases Table 23-Apr-2019 16:54 No Alerts ✓ DATA

Databases: 1

Database	Connected	Expired	Alert Level	Alert Count	Instances	Resp Time ms	Space Used %	Used Space	Free Space
ora:Standard11gR2	✔		✔		1	83.0	85.08	3,021.25	542.8

## Databases Heatmap - HTML

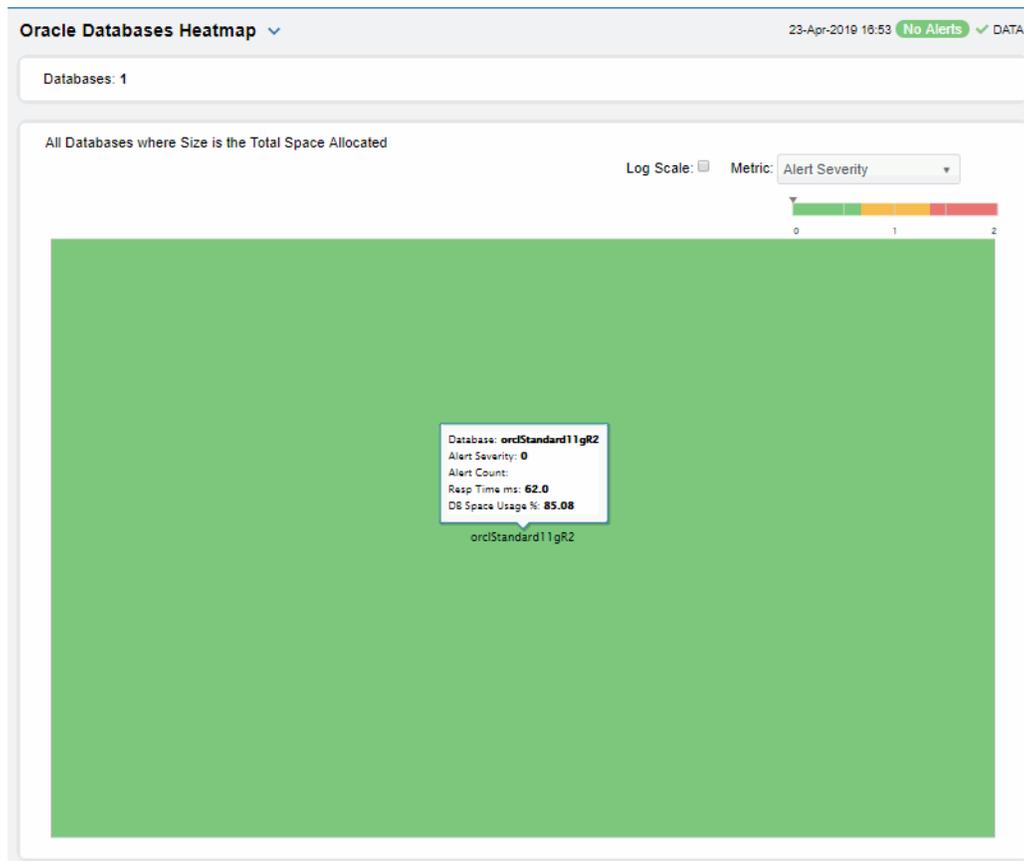
View status and alerts for all Oracle databases. Use this display to quickly identify a database with performance or utilization issues.

Choose a **Metric** from the drop-down menu: **Alert Severity**, **Alert Count**, **Response Time** or **DB Space Usage %**.

Each heatmap rectangle represents a different database. The rectangle color indicates the most critical alert state for the selected metric. Click rectangle to view details in the ["Databases Summary - HTML"](#) display and view metrics for a particular database.

Toggle between the commonly accessed Table and Heatmap displays by clicking the drop down list on the display title. Mouse-over rectangles to view more details about database performance and status.

You can view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.



### Metric:

Choose the type of metric to show in the heatmap. Each metric has its own gradient bar that maps current relative values to colors:

#### Alert Severity

The maximum level of alerts in the heatmap rectangle. Values range from **0** - **2**, as indicated in the color gradient bar, where **2** is the highest Alert Severity.

Red indicates that one or more metrics have reached their alarm threshold. Metrics that have exceeded their specified ALARM LEVEL threshold have an Alert Severity value of **2**.

Yellow indicates that one or more metrics have reached their alarm threshold. Metrics that have exceeded their specified WARNING LEVEL threshold have an Alert Severity value of **1**.

Green indicates that no metrics have reached their alert thresholds. Metrics that have not exceeded their specified thresholds have an Alert Severity value of **0**.

#### Alert Count

The total number of critical and warning alerts in the heatmap rectangle. The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.

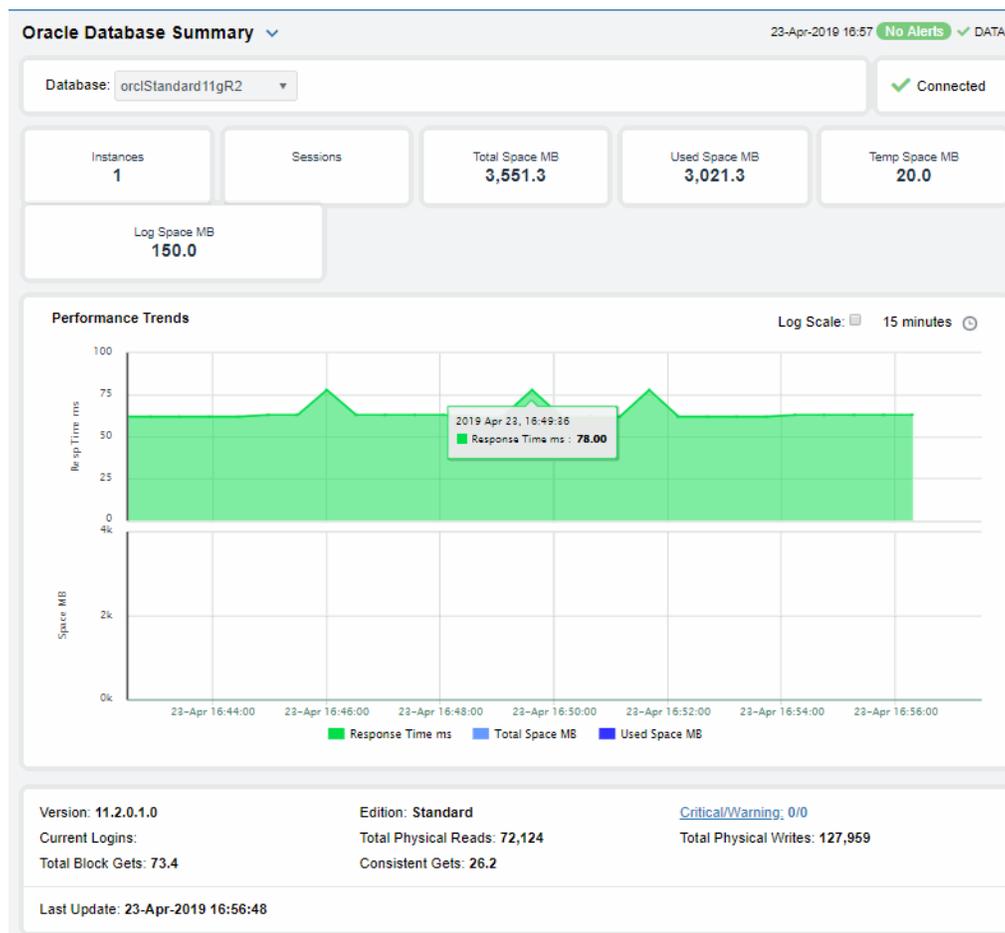
- Response Time** The amount of time, in milliseconds, since the last execution of the database. The color gradient  bar numerical values range from **0** to the maximum amount of time in the heatmap. The middle value in the gradient bar indicates the average amount.
- Database Space Usage** The amount of space used, in megabytes, in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum amount of space used in the heatmap. The middle value in the gradient bar indicates the average amount.

## Databases Summary - HTML

Track utilization and performance metrics for a specific database. Choose a **Database** from the drop-down menu. Mouse-over the utilization information boxes at the top of the display to see more details. Clicking on them takes you to the "[Oracle Databases Table - HTML](#)" display, where you can compare metrics with other Oracle databases.

The trend graph traces **Response Time**, **Total Space** and **Used Space**. Mouse-over to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.

Clicking the **Critical/Warning** link at the bottom of the display opens the Alerts Table by Component display.



## Oracle Instances View - HTML

These displays present performance metrics for Oracle Database instances. Displays are:

- ["Oracle Instances Table - HTML"](#): A list of all Oracle Database instances with detailed utilization metrics.
- ["Instances Heatmap - HTML"](#): A heatmap shows alert status of Oracle Database instances.
- ["Instance Summary - HTML"](#): Detailed performance metrics for one Oracle Database instance.

### Oracle Instances Table - HTML

Investigate and compare detailed utilization and performance metrics as well as configurations for all Oracle Database instances. This display contains all information available for Oracle Database instances such as **Instance ID**, **Instance Name**, **Expired**, **Host Name**, **Instance Role**, **Database Status**, **Version**, **SQL Hit Ratio** and **DD Hit Ratio**.

Choose one or **All** databases from the **Database** drop-down menu. Each row in the table contains data for a particular Oracle Database instance. Click a column header to sort column data in ascending or descending order. Double-click on a table row to drill-down to the ["Instance Summary - HTML"](#) display and view metrics for that particular Oracle Database instance. Toggle between the commonly accessed Heatmap and Summary displays by clicking the drop down list on the display title.

Right-click on a table cell to **Export to Excel** or **Copy Cell Value**.

The **Alert Level** column indicates the most critical alert on the instance, where:

- Red indicates that one or more alerts exceeded their ALARM LEVEL threshold in the table row.
- Yellow indicates that one or more alerts exceeded their WARNING LEVEL threshold in the table row.
- Green indicates that no alerts have reached their alert thresholds.

Oracle Instances Table 24-Apr-2019 09:19 No Alerts ✓ DATA

Database: - All - More Columns

Instances: 1

Database	Instance Id	Instance Name	Expired	Alert Level	Alert Count	Instance Role	DB Status	Version	Active Sessions	Ina
orclStandard11gR2	0	orcl		✓		PRIMARY_INSTAN	ACTIVE	11.2.0.1.0	0.0	

## Instances Heatmap - HTML

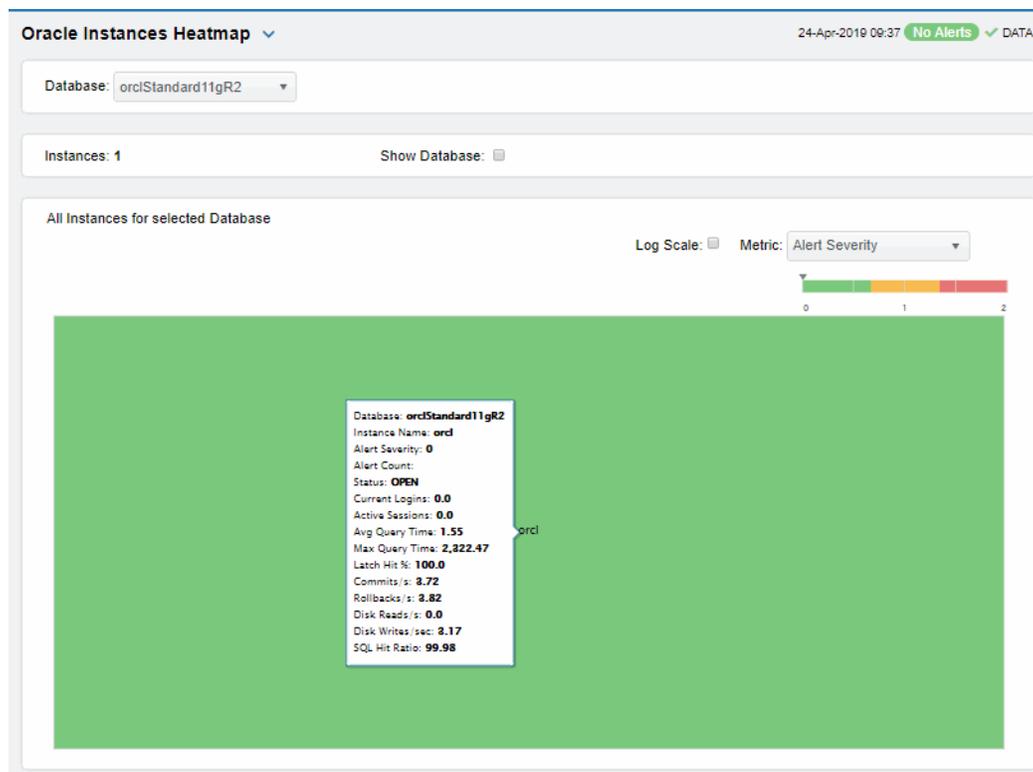
View status and alerts for all instances on one or all Oracle databases. Use this display to quickly identify an instance with performance or utilization issues.

Choose a **Metric** from the drop-down menu: **Alert Severity**, **Alert Count**, **Response Time** or **DB Space Usage %**, **Current Logins**, **Active Sessions**, **Avg Query Time**, **Max Query Time**, **Latch Hit Ratio**, **Data Dict Hit Ratio**, **SQL Hit Ratio**, **Commits/sec**, **Rollbacks/sec**, **Disk Reads/sec** or **Disk Writes/sec**.

Each heatmap rectangle represents a different instance. The rectangle color indicates the most critical alert state for the selected metric. Click rectangle to view details in the ["Instance Summary - HTML"](#) display and view metrics for a particular instance.

Toggle between the commonly accessed Table and Summary displays by clicking the drop down list on the display title. Mouse-over rectangles to view more details about database instance performance and status.

You can view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.

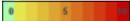


**Database** Choose a database to display.

**Instance Count** The number of instances in the display.

**Metric:**

Choose the type of metric to show in the heatmap. Each rectangle is an instance. Each metric has its own gradient bar that maps current relative values to colors:

<b>Alert Severity</b>	<p>The maximum level of alerts in the heatmap rectangle. Values range from <b>0</b> - <b>2</b>, as indicated in the color gradient  bar, where <b>2</b> is the highest Alert Severity.</p> <ul style="list-style-type: none"> <li> Red indicates that one or more metrics have reached their alarm threshold. Metrics that have exceeded their specified ALARM LEVEL threshold have an Alert Severity value of <b>2</b>.</li> <li> Yellow indicates that one or more metrics have reached their alarm threshold. Metrics that have exceeded their specified WARNING LEVEL threshold have an Alert Severity value of <b>1</b>.</li> <li> Green indicates that no metrics have reached their alert thresholds. Metrics that have not exceeded their specified thresholds have an Alert Severity value of <b>0</b>.</li> </ul>
<b>Current Logins</b>	<p>The number of users logged on. The color gradient  bar values range from <b>0</b> to the maximum count in the heatmap. The middle value in the gradient bar indicates the average count.</p>
<b>Active Sessions</b>	<p>The number of active sessions. The color gradient  bar values range from <b>0</b> to the maximum count in the heatmap. The middle value in the gradient bar indicates the average count.</p>
<b>Avg Query Time</b>	<p>The average amount of time, in seconds, to perform a query. The color gradient  bar values range from <b>0</b> to the maximum average in the heatmap. The middle value in the gradient bar indicates the average amount.</p>
<b>Max Query Time</b>	<p>The maximum amount of time, in seconds, to perform a query. The color gradient  bar values range from <b>0</b> to the maximum amount in the heatmap. The middle value in the gradient bar indicates the average maximum amount.</p>
<b>Latch Hit Ratio</b>	<p>The ratio of the number of latch misses to the number of latch gets. The color gradient  bar values range from the lowest count to the maximum count in the heatmap. The middle value in the gradient bar indicates the average.</p>
<b>Data Dict Hit Ratio</b>	<p>The ratio of logical reads to physical disk reads. The color gradient  bar values range from the lowest count to the maximum value in the heatmap. The middle value in the gradient bar indicates the average.</p>
<b>SQL Hit Ratio</b>	<p>The ratio of logical reads to physical disk reads. The color gradient  bar values range from the lowest count to the maximum value in the heatmap. The middle value in the gradient bar indicates the average.</p>
<b>Commits/sec</b>	<p>The number of commits per second. The color gradient  bar values range from <b>0</b> to the maximum count in the heatmap. The middle value in the gradient bar indicates the average.</p>
<b>Rollbacks/sec</b>	<p>The number of rollbacks per second. The color gradient  bar values range from <b>0</b> to the maximum count in the heatmap. The middle value in the gradient bar indicates the average.</p>
<b>Disk Reads/sec</b>	<p>The number of disk reads per second. The color gradient  bar values range from <b>0</b> to the maximum count in the heatmap. The middle value in the gradient bar indicates the average.</p>
<b>Disk Writes/sec</b>	<p>The number of disk writes per second. The color gradient  bar values range from <b>0</b> to the maximum count in the heatmap. The middle value in the gradient bar indicates the average.</p>

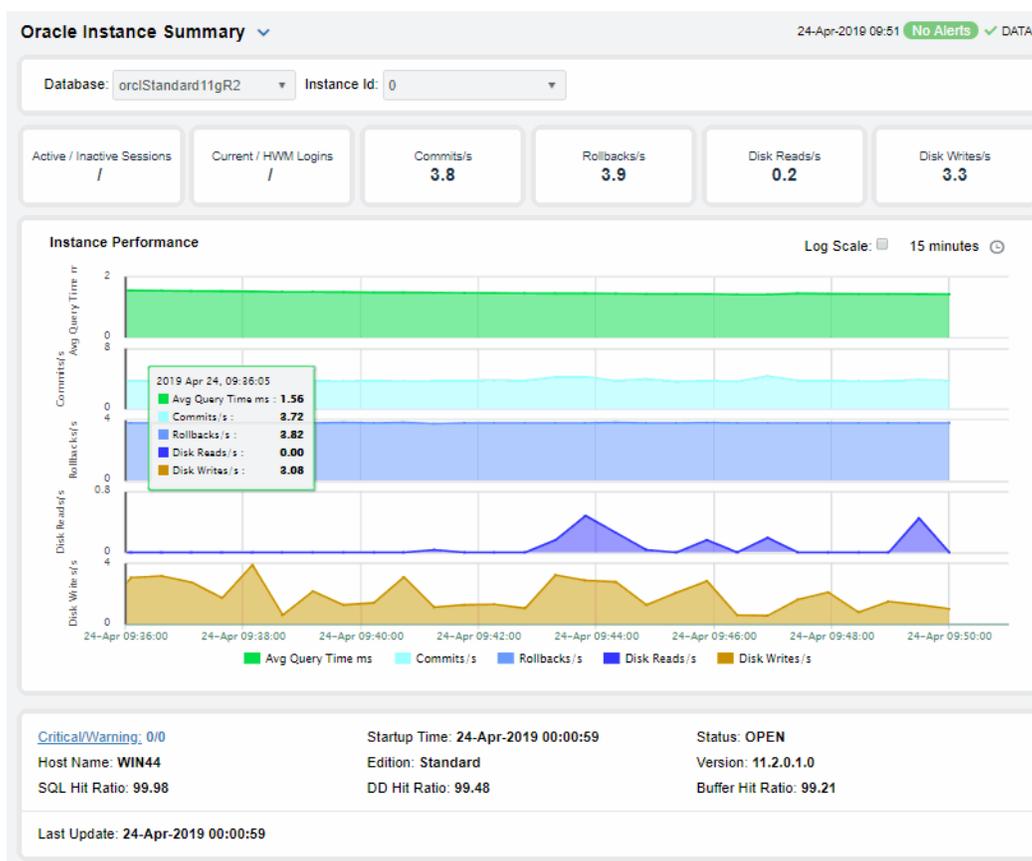
## Instance Summary - HTML

Track utilization and performance metrics for a specific instance. Choose a **Database** and an **Instance** from the drop-down menus. Mouse-over the utilization information boxes at the top of the display to see more details. Clicking on them takes you to the "Oracle Instances Table - HTML" display, where you can compare metrics with other Oracle Database instances.

The trend graph traces **Avg Query Time** in milliseconds and rates for **Commits**, **Rollbacks**, **Disk Reads** and **Disk Writes**.

Mouse-over to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.

Clicking the **Critical/Warning** link at the bottom of the display opens the Alerts Table by Component display.



---

## Oracle Database

Use the Solution Package for Oracle® Database to monitor the health of your Oracle databases. The following Oracle Database Views can be found under **Components** tab > **Databases > Oracle Database**:

- ["Database Instances View"](#)
- ["Database Details View"](#)

### Database Instances View

These displays present performance metrics and alert statuses for all Oracle databases and instances. Displays are:

- ["All Databases Heatmap"](#): Heatmap shows alert status for all Oracle databases.
- ["All Databases Table"](#): List of all Oracle databases with detailed utilization metrics.
- ["Database Summary"](#): Detailed utilization metrics and configurations for a single Oracle database.
- ["All Instances Heatmap"](#): Heatmap shows alert status of all instances on a single Oracle database.
- ["All Instances Table"](#): List of all instances on a single Oracle database with detailed utilization metrics.
- ["Instance Summary"](#): Detailed utilization metrics and configurations for a single Oracle instance.

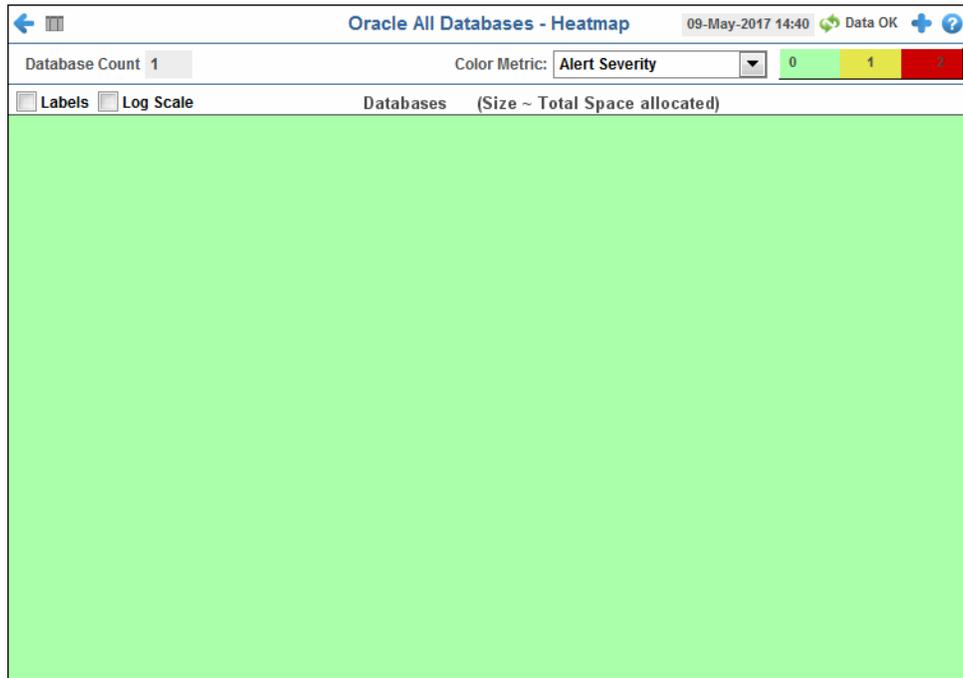
### All Databases Heatmap

View status and alerts for all Oracle databases. Use this display to quickly identify a database with performance or utilization issues.

Each heatmap rectangle represents a different database. The rectangle color indicates the most critical alert state for the selected metric.

Choose a **Color Metric** from the drop-down menu: **Alert Severity**, **Alert Count**, **Response Time** or **Database Space Usage**. Use the check-boxes  to include labels and apply log scale in the heatmap. Move your mouse over a rectangle to see additional information. Toggle between the commonly accessed **Table** and this **Heatmap** display by clicking the icon in the upper left-hand corner.

Investigate a database by clicking a heatmap rectangle to view details in the “[Database Summary](#)” display.



**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

**Database Count** The number of databases in the display.

**Metric:** Choose the type of metric to show in the heatmap. Each metric has its own gradient bar that maps current relative values to colors:

<b>Alert Severity</b>	<p>The maximum level of alerts in the heatmap rectangle. Values range from <b>0</b> - <b>2</b>, as indicated in the color gradient  bar, where <b>2</b> is the highest Alert Severity.</p> <ul style="list-style-type: none"> <li> Red indicates that one or more metrics have reached their alarm threshold. Metrics that have exceeded their specified ALARM LEVEL threshold have an Alert Severity value of <b>2</b>.</li> <li> Yellow indicates that one or more metrics have reached their alarm threshold. Metrics that have exceeded their specified WARNING LEVEL threshold have an Alert Severity value of <b>1</b>.</li> <li> Green indicates that no metrics have reached their alert thresholds. Metrics that have not exceeded their specified thresholds have an Alert Severity value of <b>0</b>.</li> </ul>
<b>Alert Count</b>	<p>The total number of critical and warning alerts in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.</p>
<b>Response Time</b>	<p>The amount of time, in milliseconds, since the last execution of the database. The color gradient  bar numerical values range from <b>0</b> to the maximum amount of time in the heatmap. The middle value in the gradient bar indicates the average amount.</p>
<b>Database Space Usage</b>	<p>The amount of space used, in megabytes, in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum amount of space used in the heatmap. The middle value in the gradient bar indicates the average amount.</p>
<b>Labels</b>	Select to include labels in the heatmap.
<b>Log Scale</b>	Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.

## All Databases Table

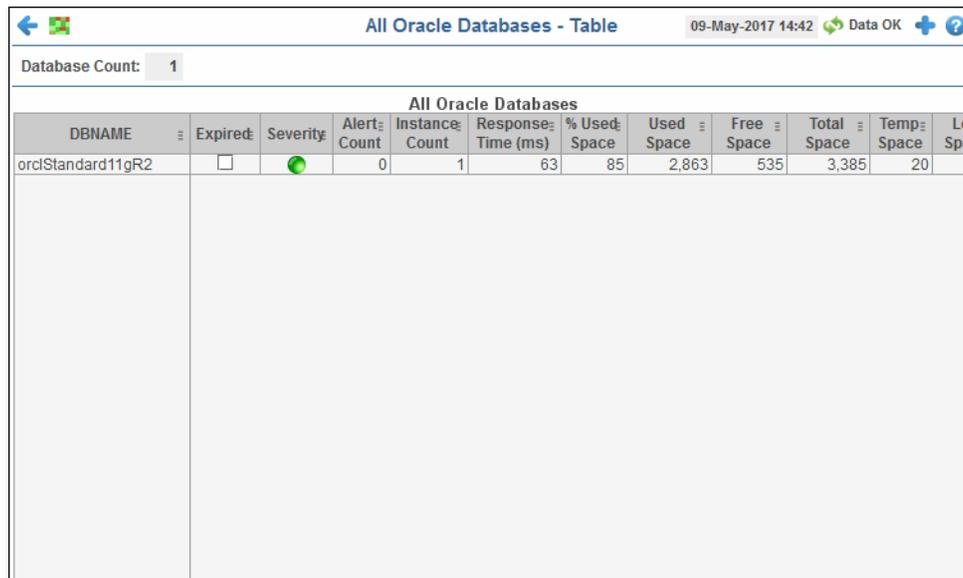
View a list of all your Oracle databases with data from the ["All Databases Heatmap"](#) display (**Alert Severity**, **Alert Count**, **Response Time**, **Database Space Usage**) in a tabular format. Each row in the table is a different Oracle database. All values refer to the database except where noted.

Light red rows indicate the database is expired or not connected.

The **Severity** column indicates the most critical alert on the database, where:

-  Red indicates that one or more alerts exceeded their ALARM LEVEL threshold in the table row.
-  Yellow indicates that one or more alerts exceeded their WARNING LEVEL threshold in the table row.
-  Green indicates that no alerts have reached their alert thresholds.

Investigate a database by clicking a row to view details in the “Database Summary” display.



All Oracle Databases												
DBNAME	Expired	Severity	Alert Count	Instance Count	Response Time (ms)	% Used Space	Used Space	Free Space	Total Space	Temp Space	Ld Sp	
orclStandard11gR2	<input type="checkbox"/>	<span style="color: green;">●</span>	0	1	63	85	2,863	535	3,385	20		

#### Title Bar (possible features are):

-   Open the previous and upper display.
-  Open an instance of this display in a new window.
-  Open the online help page for this display.
- open commonly accessed displays.

 **Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Database Count** The number of databases in the display.

#### All Oracle Databases Table

Each row is a different database.

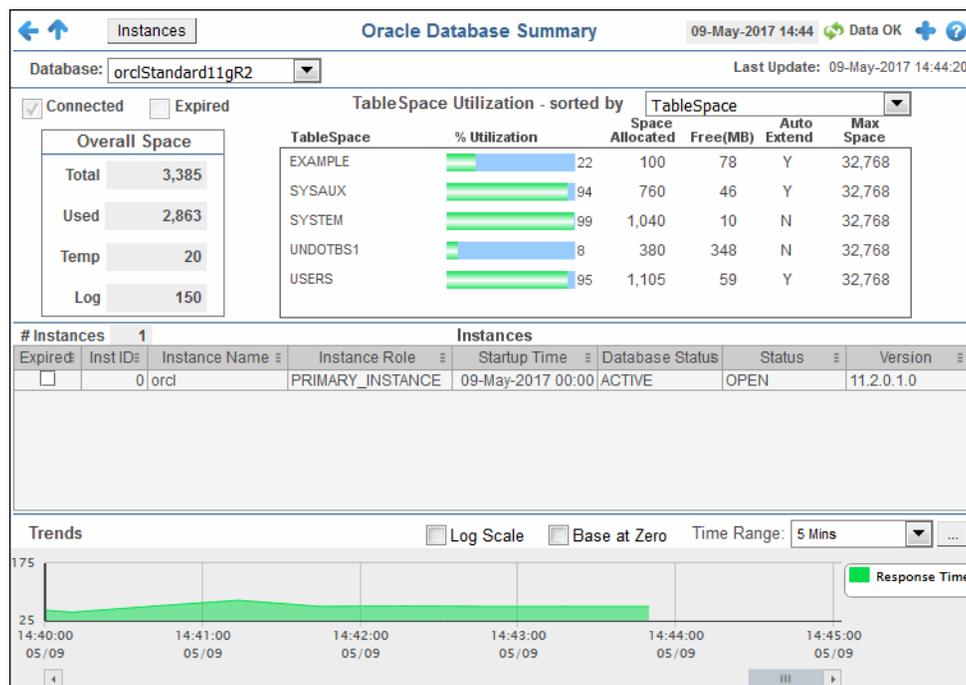
<b>DBName</b>	The database name.
<b>Expired</b>	When checked, the database is expired due to a connection issue.
<b>Severity</b>	The maximum level of alerts. <ul style="list-style-type: none"> <li><span style="color: red;">●</span> Red indicates that one or more metrics have reached their alarm threshold.</li> <li><span style="color: yellow;">●</span> Yellow indicates that one or more metrics have reached their alarm threshold.</li> <li><span style="color: green;">●</span> Green indicates that no metrics have reached their alert thresholds.</li> </ul>
<b>Alert Count</b>	The total number of critical and warning alerts.
<b>Instance Count</b>	The total number of Oracle Database instances.
<b>Response Time (ms)</b>	The amount of time, in milliseconds, since the last execution of the database.
<b>% Used Space</b>	The percent space used.
<b>Used Space</b>	The space used by the catabase, in megabytes.

<b>Free Space</b>	The amount of available space, in megabytes.
<b>Total Space</b>	The total amount of used and available space, in megabytes.
<b>Temp Space</b>	The amount of temporary space, in megabytes.
<b>Log Space</b>	The amount of log space, in megabytes.
<b>Timestamp</b>	The date and time of the last data update.

## Database Summary

View detailed performance data for a single Oracle Database, including tablespace utilization, instances on the database and response time trends. Use this display to investigate database health.

Choose a **Database** and select one metric from the **TableSpace Utilization** drop-down menu to order the tablespace rows by that metric. The order is ascending. Mouse-over the trend graph to see additional information. Investigate an instance by clicking a row in the **Instances** table to view details in the "Instance Summary" display.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

<b>Database</b>	Choose a database to display.
<b>Connected</b>	When checked, the database is connected.
<b>Expired</b>	When checked, the database is expired due to a connection issue.
<b>Overall Space</b>	Values refer to the selected database.
<b>Total</b>	The total amount of space, in megabytes.
<b>Used</b>	The amount of space used, in megabytes.
<b>Temp</b>	The total amount of temporary space, in megabytes.
<b>Log</b>	The total amount of log space, in megabytes.
<b>TableSpace Utilization</b>	Choose a tablespace utilization metric to show data for (in the box below the drop-down menu). Values refer to the tablespaces on the selected database: <ul style="list-style-type: none"> <li>• <b>TableSpace:</b> Tablespace names in alphabetical order.</li> <li>• <b>% Utilization:</b> The percent utilization of each tablespace in ascending order.</li> <li>• <b>Allocated:</b> The allocated space for each tablespace in ascending order.</li> </ul>
<b>TableSpace</b>	The tablespace name.
<b>% Utilization</b>	The percent utilization.
<b>Space Allocated</b>	The amount of allocated space, in megabytes.
<b>Free(MB)</b>	The amount of free space, in megabytes.
<b>Auto Extend</b>	Indicates whether auto extend is enabled (Y/N).
<b>Max Space</b>	The maximum amount of allocated space for the tablespace, in megabytes.
<b># Instances</b>	Each table row is an instance on the selected database. Click an instance to see details in the <a href="#">"Instance Summary"</a> display.
<b>Expired</b>	When checked, the instance is expired due to a connection issue.
<b>Inst ID</b>	The unique identifier for the instance.
<b>Instance Name</b>	The name of the instance.
<b>Instance Role</b>	The role of the instance.
<b>Startup Time</b>	The date and time the instance began.
<b>Database Status</b>	The status of the database this instance is on.
<b>Status</b>	The status of this instance.
<b>Version</b>	The version number.

**Trends**

Traces response time of the selected database.

**Log Scale**

Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.

**Base at Zero**

Use zero as the Y axis minimum for all graph traces.

**Time Range**

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar  .

By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

**All Instances Heatmap**

View status and alerts of all instances on an Oracle Database. Use this display to quickly identify instances with performance or utilization issues.

Each heatmap rectangle represents a different instance. The rectangle color indicates the most critical alert state for the selected metric.

Choose a **Color Metric** from the drop-down menu, such as **Alert Severity**, **Active Sessions** and **Disk Reads**. Use the check-boxes  to include or exclude labels and use log scale in the heatmap. Move your mouse over a rectangle to see additional information. Toggle between the commonly accessed **Table** and this **Heatmap** display by clicking the icon in the upper left-hand corner. Mouse-over rectangles to view more instance performance details.

Investigate an instance by clicking a heatmap rectangle to view details in the “[Instance Summary](#)” display.



**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

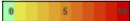
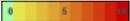
**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Database** Choose a database to display.

**Instance Count** The number of instances in the display.

**Metric:**

Choose the type of metric to show in the heatmap. Each rectangle is an instance. Each metric has its own gradient bar that maps current relative values to colors:

<b>Alert Severity</b>	<p>The maximum level of alerts in the heatmap rectangle. Values range from <b>0</b> - <b>2</b>, as indicated in the color gradient  bar, where <b>2</b> is the highest Alert Severity.</p> <ul style="list-style-type: none"> <li> Red indicates that one or more metrics have reached their alarm threshold. Metrics that have exceeded their specified ALARM LEVEL threshold have an Alert Severity value of <b>2</b>.</li> <li> Yellow indicates that one or more metrics have reached their alarm threshold. Metrics that have exceeded their specified WARNING LEVEL threshold have an Alert Severity value of <b>1</b>.</li> <li> Green indicates that no metrics have reached their alert thresholds. Metrics that have not exceeded their specified thresholds have an Alert Severity value of <b>0</b>.</li> </ul>
<b>Current Logins</b>	The number of users logged on. The color gradient  bar values range from <b>0</b> to the maximum count in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Active Sessions</b>	The number of active sessions. The color gradient  bar values range from <b>0</b> to the maximum count in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Avg Query Time</b>	The average amount of time, in seconds, to perform a query. The color gradient  bar values range from <b>0</b> to the maximum average in the heatmap. The middle value in the gradient bar indicates the average amount.
<b>Max Query Time</b>	The maximum amount of time, in seconds, to perform a query. The color gradient  bar values range from <b>0</b> to the maximum amount in the heatmap. The middle value in the gradient bar indicates the average maximum amount.
<b>Latch Hit Ratio</b>	The ratio of the number of latch misses to the number of latch gets. The color gradient  bar values range from the lowest count to the maximum count in the heatmap. The middle value in the gradient bar indicates the average.
<b>Data Dict Hit Ratio</b>	The ratio of logical reads to physical disk reads. The color gradient  bar values range from the lowest count to the maximum value in the heatmap. The middle value in the gradient bar indicates the average.
<b>SQL Hit Ratio</b>	The ratio of logical reads to physical disk reads. The color gradient  bar values range from the lowest count to the maximum value in the heatmap. The middle value in the gradient bar indicates the average.
<b>Commits/sec</b>	The number of commits per second. The color gradient  bar values range from <b>0</b> to the maximum count in the heatmap. The middle value in the gradient bar indicates the average.
<b>Rollbacks/sec</b>	The number of rollbacks per second. The color gradient  bar values range from <b>0</b> to the maximum count in the heatmap. The middle value in the gradient bar indicates the average.
<b>Disk Reads/sec</b>	The number of disk reads per second. The color gradient  bar values range from <b>0</b> to the maximum count in the heatmap. The middle value in the gradient bar indicates the average.
<b>Disk Writes/sec</b>	The number of disk writes per second. The color gradient  bar values range from <b>0</b> to the maximum count in the heatmap. The middle value in the gradient bar indicates the average.
<b>Labels</b>	Select to include labels in the heatmap.
<b>Log Scale</b>	Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.

## All Instances Table

View data from the “All Instances Heatmap” display in a tabular format, as well as additional metrics and configuration information about instances.

Choose a database from the drop-down menu. Each row in the table is a different Oracle Database instance on the selected database. Investigate an instance by clicking a row to view details in the “Instance Summary” display.

Light red rows indicate the following:

- The database has an issue (the **Database Status** value is not **ACTIVE** or the **Active State** value is not **NORMAL**)
- The instance has an issue (the instance expired or the **Instance Status** value is not **OPEN**)

DBNAME	Instance ID	Instance Name	Expired	Severity	Host Name	Instance Role	Database Status	Instance Status
orclStandard11gR2	0	orcl	<input type="checkbox"/>	●	WIN44	PRIMARY_INSTANCE	ACTIVE	OPEN

### Title Bar (possible features are):

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu, Table open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

**Database** Choose a database to display.

**Instance Count** The number of instances in the display.

**Oracle Databases Instances Table**

Each row is a different instance. Column values refer to the instance except where noted. This data is obtained from the vendor. See vendor documentation for details.

<b>DBName</b>	The name of the database the instance is using.
<b>Instance ID</b>	The unique identifier for the instance.
<b>Instance Name</b>	The instance name.
<b>Expired</b>	When checked, the database is expired due to a connection issue.
<b>Severity</b>	The maximum level of alerts.  Red indicates that one or more metrics have reached their alarm threshold.  Yellow indicates that one or more metrics have reached their alarm threshold.  Green indicates that no metrics have reached their alert thresholds.
<b>Host Name</b>	The host for the instance.
<b>Instance Role</b>	The instance role. Values are: <ul style="list-style-type: none"> <li>• PRIMARY_INSTANCE</li> <li>• SECONDARY_INSTANCE</li> <li>• UNKNOWN</li> </ul>
<b>Database Status</b>	The current database status. Values are: <ul style="list-style-type: none"> <li>• ACTIVE</li> <li>• SUSPENDED</li> <li>• INSTANCE_RECOVERY</li> </ul>
<b>Instance Status</b>	The current state of the instance. Values are: <ul style="list-style-type: none"> <li>• STARTED</li> <li>• MOUNTED</li> <li>• OPEN</li> <li>• OPEN MIGRATE</li> </ul>
<b>Active State</b>	<ul style="list-style-type: none"> <li>• The database state when active. Values are:</li> <li>• NORMAL</li> <li>• QUIESCING</li> <li>• QUIESCED</li> </ul>
<b>Archiver</b>	This data is obtained from the vendor. See vendor documentation for details.
<b>Blocked</b>	Is the instance in block mode? <b>Yes/No</b>
<b>Parallel</b>	Is the instance in parallel mode? <b>Yes/No</b>
<b>Shutdown Pending</b>	Is a shutdown pending? <b>Yes/No</b>
<b>Start Time</b>	The date and time the instance started.
<b>Edition</b>	The edition number of the instance.
<b>Version</b>	The instance software version.

<b>Host OS</b>	The host operating system.
<b>CPU Count</b>	The database CPU count.
<b>Avg Query Time</b>	The average amount of time, in seconds, to perform a query.
<b>Min Query Time</b>	The minimum amount of time, in seconds, to perform a query.
<b>Max Query Time</b>	The maximum amount of time, in seconds, to perform a query.
<b>Buffer Hit Ratio</b>	The ratio of the number of latch misses to the number of latch gets.
<b>DD Hit Ratio</b>	The ratio of logical reads to physical disk reads.
<b>SQL Hit Ratio</b>	The ratio of logical SQL reads to physical disk reads.
<b>Logins</b>	This data is obtained from the vendor. See vendor documentation for details.
<b>Current Logins</b>	The number of users currently logged in.
<b>Cum Logins</b>	The total number of logged in users since the instance started.
<b>Logins Highwater</b>	The login limit for the database.
<b>Active Sessions</b>	The number of logins that are currently active.
<b>Inactive Sessions</b>	The number of logins that are currently inactive.
<b>Other Sessions</b>	This data is obtained from the vendor. See vendor documentation for details.
<b>Block Gets</b>	The number of gets in block mode.
<b>Consistent Gets</b>	The number of gets in consistent mode.
<b>Physical Get Reads</b>	The number of physical gets reads.
<b>Latch Hit%</b>	The amount of latch gets, in percent.
<b>Latch Miss%</b>	The amount of latch misses, in percent.
<b>Latch Spin%</b>	This data is obtained from the vendor. See vendor documentation for details.
<b>Parse Time%</b>	This data is obtained from the vendor. See vendor documentation for details.
<b>Recursive%</b>	This data is obtained from the vendor. See vendor documentation for details.
<b>Other%</b>	This data is obtained from the vendor. See vendor documentation for details.
<b>Deadlocks</b>	The number of deadlocks since the database started.
<b>Commits/sec</b>	The number of commits per second.
<b>Rollbacks/sec</b>	The number of rollbacks per second.
<b>Block Reads/sec</b>	The number of block reads per second.

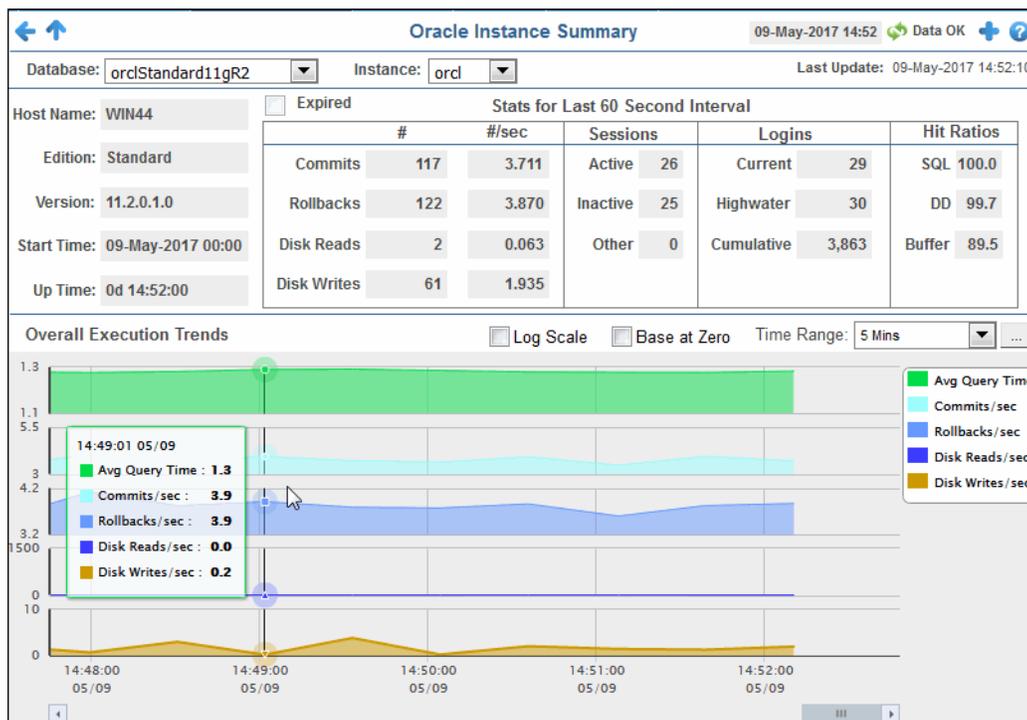
**Block Writes/sec** The number of block writes per second.

**Timestamp** The date and time of the last data update.

### Instance Summary

View detailed data for a single Oracle Database instance, such as performance statistics for last 60 seconds, overall usage and execution trends. Use this display to closely investigate the health of an instance.

Choose a **Database** and **Instance** metric from the drop-down menus. Mouse-over the trend graph to see additional information.



**Title Bar (possible features are):**

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu Table open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Database** Choose a database to display.

**Instance** Choose an instance to display.

**Host Name** The name of the host.

**Edition** This data is obtained from the vendor. See vendor documentation for details.

<b>Version</b>	This data is obtained from the vendor. See vendor documentation for details.
<b>Start Time</b>	The date and time the instance started.
<b>Up Time</b>	The instance session duration.
<b>Expired</b>	When checked, the database is expired due to a connection issue.

**Stats for Last 60 Second Interval**

Values refer to the selected instance.

	<b>#</b>	<b>#/sec</b>
<b>Commits</b>	The number of commits since the instance started.	The current number of commits per second.
<b>Rollbacks</b>	The number of rollbacks since the instance started.	The current number of rollbacks per second.
<b>Disk Reads</b>	The number of disk reads since the instance started.	The current number of disk reads per second.
<b>Disk Writes</b>	The number of disk writes since the instance started.	The current number of disk writes per second.
<b>Sessions</b>	<b>Active</b>	The current number of active sessions on the instance.
	<b>Inactive</b>	The current number of inactive sessions on the instance.
	<b>Other</b>	This data is obtained from the vendor. See vendor documentation for details.
<b>Logins</b>	<b>Current</b>	The current number of logins on the instance.
	<b>Highwater</b>	The login limit for the instance.
	<b>Cumulative</b>	This data is obtained from the vendor. See vendor documentation for details.
<b>Hit Ratios</b>	<b>SQL</b>	The ratio of logical SQL reads to physical disk reads.
	<b>DD</b>	The ratio of logical reads to physical disk reads.
	<b>Buffer</b>	The ratio of the number of latch misses to the number of latch gets.

**Overall Execution Trends**

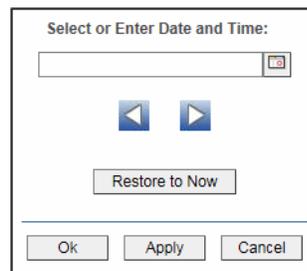
Traces the following for the selected instance:

**Avg Query Time:** The average amount of time to perform a query, in seconds.

- **Commits/sec:** The number of commits per second.
- **Rollbacks/sec:** The number of rollbacks per second.
- **Disk Reads/sec:** The number of disk reads per second.
- **Disk Writes/sec:** The number of disk writes per second.

**Log Scale** Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.

- Base at Zero** Use zero as the Y axis minimum for all graph traces.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Database Details View

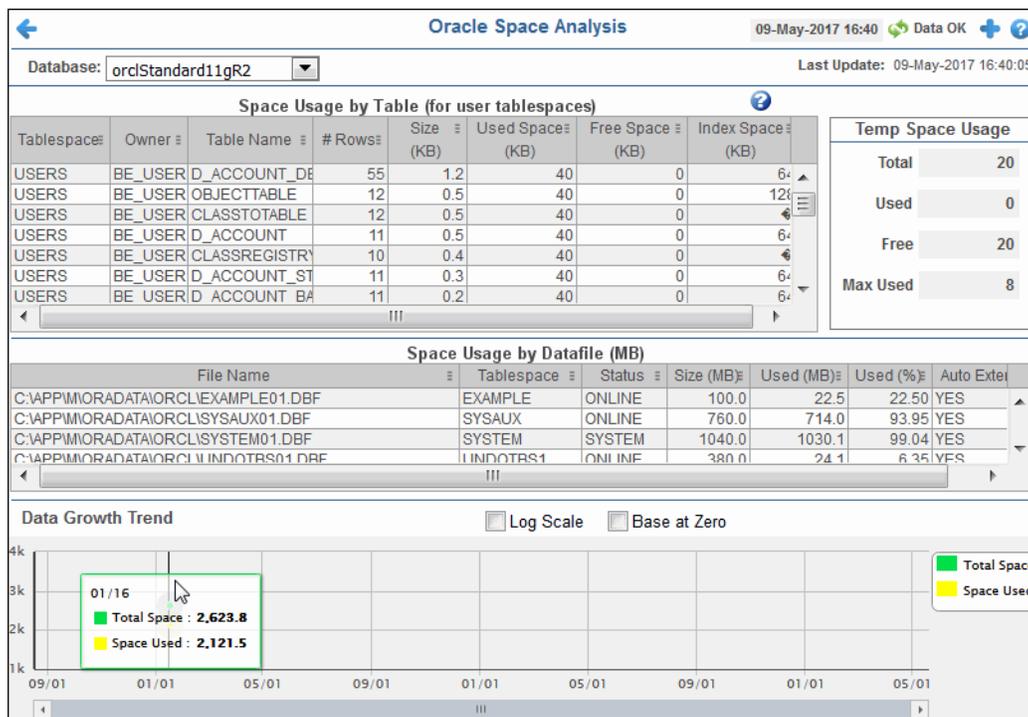
These displays present detailed database utilization metrics for an Oracle database. Displays are:

- **"Space"**: Shows detailed space utilization metrics for a single database.
- **"Contention"**: Shows detailed contention data for a single database.
- **"Processes"**: Shows detailed CPU and memory utilization per process for a single database.
- **"Sessions"**: Shows detailed CPU and memory utilization per session and user on a single database.
- **"Queries"**: Shows detailed query performance data such as disk utilization, execution and wait times and full SQL texts.
- **"Reports"**: Shows the following detailed reports: Audited Objects, Audited Privileges, Database Initialization, Feature Usage and Invalid Objects.

## Space

Analyse tablespace utilization and trends at the data file level. Choose a database from the drop-down menu to view table and data file utilization data in a tabular format. Each row in the upper table is a different tablespace on the selected database. Each row in the lower table is a different data file on the selected database.

Mouseover the trend graph to see metrics for a specific day or time.



**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Database** Choose a database to display.

**Space Usage by Table (for user tablespaces)**

Each row is a different table on the selected database. Column values refer to the table.

- Tablespace** The tablespace in which the table resides.
- Owner** The table owner.
- Table** The table name.
- #Rows** The current number of rows in the table.
- Size (KB)** The table size, in kilobytes.
- Used Space (KB)** The amount of space used by the table, in kilobytes.
- Free Space (KB)** The amount of free space for the table, in kilobytes.

**Index Space (KB)** The amount of index space for the table, in kilobytes.

#### Temp Space Usage

**Total** The total amount of temp space for the table, in kilobytes.

**Used** The amount of used space for the table, in kilobytes.

**Free** The amount of free temp space for the table, in kilobytes.

**Max Used** The maximum amount of space used by the table, in kilobytes.

#### Space Usage by Datafile (MB)

**File Name** The name of the data file.

**Tablespace** The name of the tablespace in which the data file resides.

**Status** The data file status:

- ONLINE
- SYSTEM

**Size (MB)** The data file size, in megabytes.

**Used (MB)** The amount of space used, in megabytes.

**Auto Extend** Describes whether auto extend is configured (Yes/No).

#### Data Growth Trend

Traces the following for the selected database:

- **Total Space:** The total amount of space for the table, in kilobytes.
- **Space Used:** The total amount of space used for the table, in kilobytes.

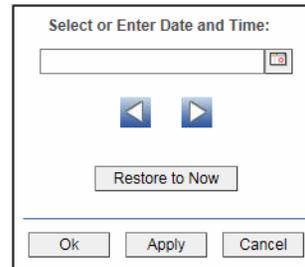
**Log Scale** Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.

**Base at Zero**

Use zero as the Y axis minimum for all graph traces.

**Time Range**

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

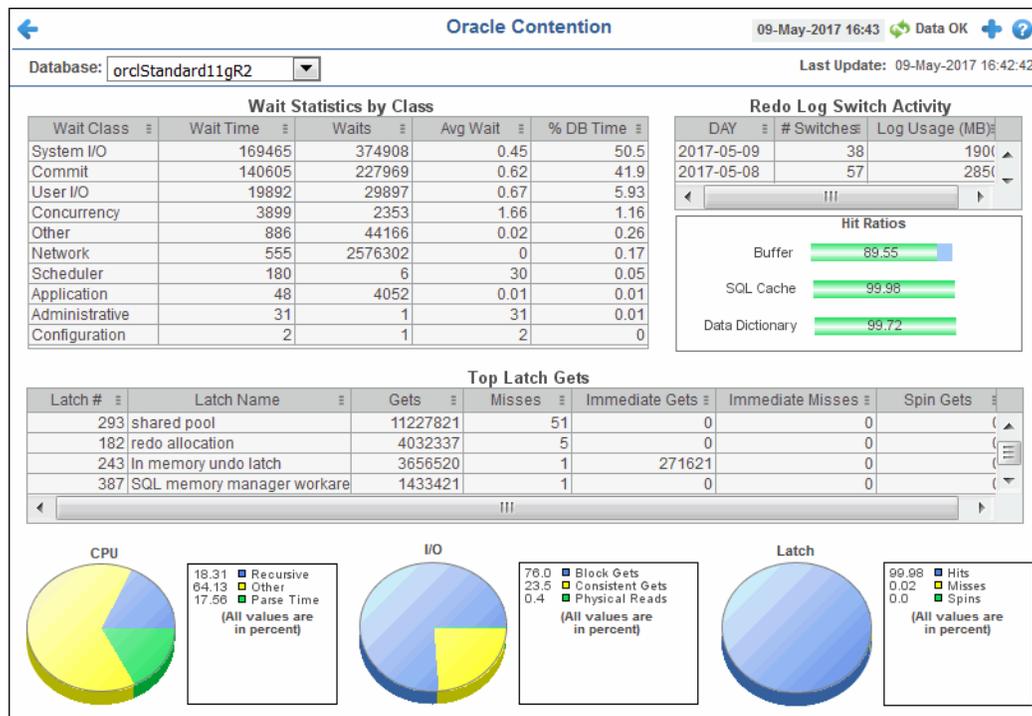
Click **Restore to Now** to reset the time range end point to the current time.

**Contention**

Analyse contention issues impacting database performance, including wait statistics per class, log switch activity, hit ratios and top latch gets.

Choose a database from the drop-down menu. Each row in the upper table is a different wait class on the selected database. Each row in the lower table is a different top latch on the selected database.

Mouseover the pie graph to see more information about database utilization.



**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Database** Choose a database to display.

**Wait Statistics by Class**

Each row is a different wait class on the selected database. Column values refer to the wait class.

- Wait Class** The wait class name.
- Wait Time** The amount of time, in milliseconds, for the class to process an event.
- Waits** The number of wait events.
- Avg Wait** The average amount of time, in seconds, for the class to process an event.
- % DB Time** The amThis data is obtained from the vendor. See vendor documentation for details.

**Redo Log Switch Activity**

Each row is a different day. Column values refer to the selected database.

- DAY** The date of the data update.

- #Switches**                    The number of switch processes performed by the database for the day.
- Log Usage (MB)**            This data is obtained from the vendor. See vendor documentation for details.

**Hit Ratios**

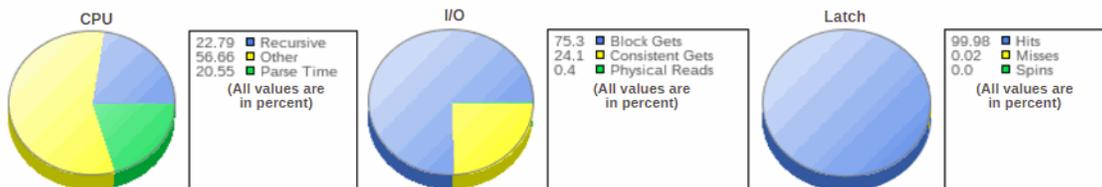
Values refer to the selected database.

- Buffer**                        The ratio of the number of latch misses to the number of latch gets.
- SQL Cache**                 The ratio of logical SQL reads to physical disk reads.
- Data Dictionary**            The ratio of logical reads to physical disk reads.

**Top Latch Gets**

Each row is a different latch on the selected database. Column values refer to the latch.

- Latch #**                        The unique identifier for the latch.
- Latch Name**                 The latch name.
- Gets**                            The number of gets for the latch.
- Misses**                        The number of misses for the latch.
- Immediate Gets**            The number of immediate gets for the latch.
- Immediate Misses**         The number of immediate misses for the latch.
- Spin Gets**                    The number of spin gets for the latch.



**CPU**

Pie graph represents the selected database. Values are the percent utilization of **Recursive** calls, **Parse** time, and **Other** background processes (such as looking for buffers and fetching).

**I/O**

Pie graph represent the selected database. Values are the percent utilization of **Block Gets** calls, **Consistent Gets** and **Physical Reads**.

**Latch**

Pie graph represents the selected database. Values are the percent utilization of **Hits**, **Misses** and **Spins**.

**Processes**

Find out which persons and processes are using the most CPU on a database.

Choose a database from the drop-down menu. Each row in the upper table is a different user on the selected database. Each row in the lower table is a different process on the selected database.

The screenshot shows the Oracle Process Details window for the database 'orclStandard11gR2'. It features a title bar with navigation icons, a 'Data OK' status indicator, and a timestamp of '09-May-2017 16:53'. The 'Last Update' is '09-May-2017 16:52:37'. Two tables are displayed:

**Top CPU Hogs**

User	SID	CPU (sec)	CPU Recursive	CPU Parsetime
SYSTEM	42	121.28	2.73	2
RTVHISTORY	17	120.33	74.52	90
SYSTEM	28	80.53	0.23	0
SYSTEM	55	79.97	0.05	0

**Process Memory Usage**

PID	SPID	Program	Max Memory Used	Memory Used	Freeable Memory
10	2716	ORACLE.EXE (DBW0)	13.04	13.04	6.15
51	5796	ORACLE.EXE (SHAD)	29.83	10.45	1.4
37	5316	ORACLE.EXE (SHAD)	40.95	9.95	0.7
21	1288	ORACLE.EXE (CJQ0)	5.84	5.84	4.1
11	2720	ORACLE.EXE (LGWR)	5.59	5.59	0.1
13	2728	ORACLE.EXE (SMON)	5.27	5.27	3.9
15	2736	ORACLE.EXE (MMON)	4.26	4.26	2.1
27	3096	ORACLE.EXE (Q002)	4.25	4.25	0.1
24	5288	ORACLE.EXE (SHAD)	38.14	4.02	0.8
40	5556	ORACLE.EXE (SHAD)	6.89	3.7	0.1
38	5480	ORACLE.EXE (SHAD)	47.64	2.77	0.1

**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Database** Choose a database to display.

**Top CPU Hogs**

Each row is a different user on the selected database. Column values refer to the user.

<b>User</b>	The user name.
<b>SID</b>	The unique user ID.
<b>CPU (sec)</b>	The amount of database CPU used per second, in megabytes.
<b>CPU Recursive</b>	The amount of CPU used for recursive calls per second, in megabytes.
<b>CPU Parsetime</b>	The amount of CPU used for parsetime per second, in megabytes.

**Process Memory Usage**

Each row is a different process on the selected database. Column values refer to the process.

<b>PID</b>	The unique process ID.
<b>SPID</b>	The unique SQL process ID.
<b>Program</b>	The name of the application.

- Max Memory Used**            The maximum amount of database memory used, in megabytes.
- Memory Used**                The maximum amount of database memory used, in megabytes.
- Freeable Memory**            The amount of available database memory for the process, in megabytes.

## Sessions

Analyse read/write response time, CPU utilization, locks, blocks and memory allocation per session.

The screenshot shows the 'Oracle Session Details' window for database 'orclStandard11gR2'. It contains several data tables:

User Name	OS User	PID	SID	Serial #	Physical Reads	Block Gets	Consistent Gets
RTVHISTORY	m	1234	32	42	26261	3	39131
RTVHISTORY	m	1234	40	11	10803	816360	38361
RTVHISTORY	m	1234	17	13	4588	1768223	25327
Oracle	SYSTEM	2728	13	1	4033	26779	3491

User Name	SID	CPU Usage (sec)	CPU Recursive (sec)	CPU Parsetime (sec)
SYSTEM	42	122.78	2.93	3.04
RTVHISTORY	17	121.53	75.44	91.8
SYSTEM	28	81.93	0.23	0.3
SYSTEM	55	81.35	0.05	0

User Name	SID	Lock Mode	Lock Type
Oracle	12	1	XR
Oracle	12	1	RD
Oracle	42	1	CF

User name	OS User	Machine	Program	Memory (MB)
DBSNMP	NT AUTHORITY\WORKGROUP	emagent.exe		8.9
DBSNMP	NT AUTHORITY\WORKGROUP	emagent.exe		8.7
Oracle	SYSTEM	WORKGROUP	ORACLE_EYE (DBSNMP)	8.5

Blocking SID	Blocked SID	Serial #	User Name	OS User	Machine	Wait Class	Seconds in Wait
--------------	-------------	----------	-----------	---------	---------	------------	-----------------

**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu , Table open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

- Database**    Choose a database to display.
- Top Sessions for Physical IO**  
Each row is a different session on the selected database. Column values refer to the session.
- User Name**            The user name.
- OS User**                This data is obtained from the vendor. See vendor documentation for details.

<b>PID</b>	The unique process ID.
<b>SID</b>	The unique SQL process ID.
<b>Serial #</b>	This data is obtained from the vendor. See vendor documentation for details.
<b>Physical Reads</b>	The number of physical gets reads.
<b>Block Gets</b>	The number of gets in block mode.
<b>Consistent Gets</b>	The number of gets in consistent mode.

### Top Sessions by CPU Utilization

Each row is a different session on the selected database. Column values refer to the session.

<b>User Name</b>	The name of the user.
<b>SID</b>	The unique SQL process ID.
<b>CPU Usage (sec)</b>	The amount of CPU used per second, in megabytes.
<b>CPU Recursive (sec)</b>	The amount of CPU used for recursive calls per second, in megabytes.
<b>CPU Parsetime (sec)</b>	The amount of CPU used for parsetime per second, in megabytes.

### Sessions Holding Locks

Each row is a different lock on the selected database. Column values refer to the session.

<b>User Name</b>	The name of the user with the lock.
<b>SID</b>	The unique SQL process ID.
<b>Lock Mode</b>	The lock mode.
<b>Lock Type</b>	The lock type.

### Session Memory Allocation

Each row is a different session on the selected database. Column values refer to the session.

<b>User Name</b>	The name of the user with the blocked SQL process.
<b>OS User</b>	This data is obtained from the vendor. See vendor documentation for details.
<b>Machine</b>	The unique identifier of the user's machine.
<b>Program</b>	The name of the application.
<b>Memory (MB)</b>	The amount of memory used, in megabytes.

### Blocked Sessions

Each row is a different session on the selected database. Column values refer to the session with a blocked SQL process and the SQL process causing the block.

<b>Blocking SID</b>	The process ID of the SQL process that is blocking.
<b>Blocked SID</b>	The ID of the SQL process that is blocked.
<b>Serial #</b>	This data is obtained from the vendor. See vendor documentation for details.

<b>User Name</b>	The name of the user with the blocked SQL process.
<b>OS User</b>	This data is obtained from the vendor. See vendor documentation for details.
<b>Machine</b>	The unique identifier of the user's machine.
<b>Wait Class</b>	The wait class.
<b>Seconds in Wait</b>	The length of the wait, in seconds.

## Queries

Analyse database query disk utilization, execution times and full table scans per SQL process. Choose a database from the drop-down menu, then toggle between **High Disk I/O**, **High Execution Time** or **Full Table Scans**. Click a row to populate fields below the table, such as wait time, rows processed and disk reads.

The screenshot shows the Oracle Query Analysis interface. At the top, it displays the database name 'orclStandard11gR2' and the date '09-May-2017 16:57'. Below this, there are three radio buttons for filtering queries: 'High Disk I/O' (selected), 'High Execution Time', and 'Full Table Scans'. A table lists the top queries with columns for SQL ID, SQL Text, Session ID, and Avg I/O per Execution. The selected query has SQL ID 'a5z18znkd851f' and SQL Text 'select 'orclStandard11gR2' as DBNAME, (select sum(bytes)/1024/1024 from v\$datafile) as total\_data, (select sum(bytes)/1024/1024 from dba\_free\_space) free\_data, (select sum(bytes)/1024/1024 from dba\_segments) used\_data, (select sum(bytes)/1024/1024 from v\$tempfile) as temp, (select sum(bytes)/1024/1024 from v\$log) as log from dual'. Below the table, the 'Full SQL Text' is displayed. At the bottom, there are three summary boxes: 'Summary' (Rows Processed: 2, Executions: 2, Fetches: 2, Sorts: 0, Optimizer Mode: ALL\_ROWS, Optimizer Cost: 2), 'Wait Times (seconds)' (Application: 0.102, Concurrency: 0.102, User I/O: 0.313, PL SQL Exec: 0.004, Java Exec: 0.000), and 'I/O' (Disk Reads: 55, Phys Read Requests: 251, Phys Read Bytes: 3563520, Phys Write Requests: 0, Phys Write Bytes: 0, Buffer Gets: 42334).

### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Database** Choose a database to display.

**Top Queries for:**

Each row is a different SQL process on the selected database. Column values refer to the SQL process.

**High Disk I/O**

<b>SQL ID</b>	The SQL Id.
<b>SQL Text</b>	The textual content.
<b>Session ID</b>	The unique identifier for the session.
<b>Avg I/O per Execution</b>	The average amount of data input and output per execution, in megabytes.

**High Execution Time**

<b>SQL ID</b>	The SQL Id.
<b>SQL Text</b>	The textual content.
<b>Session ID</b>	The unique identifier for the session.
<b>Execution Time Sec</b>	The amount time to execute per second, in megabytes.

**Full Table Scans**

<b>SQL ID</b>	The SQL Id.
<b>SQL Text</b>	The textual content.
<b>Operation Cost</b>	This data is obtained from the vendor. See vendor documentation for details.
<b>CPU Cost</b>	This data is obtained from the vendor. See vendor documentation for details.
<b>IO Cost</b>	This data is obtained from the vendor. See vendor documentation for details.
<b>Temp Space</b>	This data is obtained from the vendor. See vendor documentation for details.

**Full SQL Text** The SQL textual content.

**Summary**

Values refer to the SQL process selected in the table.

<b>Rows Processed</b>	The total number of rows processed since the process started.
<b>Executions</b>	The total number of executions since the process started.
<b>Fetches</b>	The total number of fetches since the process started.
<b>Sorts</b>	The total number of sorts since the process started.
<b>Optimizer Mode</b>	The optimizer mode used.

**Optimizer Cost** The total number of single block disk reads since the process started.

### Wait Times

Values refer to the SQL process selected in the table.

**Application** The name of the application.

**Concurrency** This data is obtained from the vendor. See vendor documentation for details.

**User I/O** This data is obtained from the vendor. See vendor documentation for details.

**PL SQL Exec** This data is obtained from the vendor. See vendor documentation for details.

**Java Exec** This data is obtained from the vendor. See vendor documentation for details.

### I/O

Values refer to the SQL process selected in the table.

**Disk Reads** The total number of disk reads since the process started.

**Phys Read Requests** The total number of physical disk reads since the process started.

**Phys Read Bytes** The total number of physical disk reads since the process started.

**Phys Write Requests** The total number of physical disk writes since the process started.

**Phys Write Bytes** The total number of bytes written since the process started.

**Buffer Gets** The total number of gets since the process started.

## Reports

Analyse detailed database activity. Choose a database and a report from the drop-down menus. Reports available are: **Audited Objects, Audited Privileges, Database Initialization, Feature Usage and Invalid Objects.**

Parameter	Value	Type	Description	Modified	Dynamic	Basic
active_instance_count		Integer	number of active instances in the cluster datab	No	No	No
aq_tm_processes	0	Integer	number of AQ Time Managers to start	No	Yes	No
archive_lag_target	0	Integer	Maximum number of seconds of redos the sta	No	Yes	No
asm_diskgroups		String	disk groups to mount automatically	No	Yes	No
asm_diskstring		String	disk set locations for discovery	No	Yes	No
asm_power_limit	1	Integer	number of processes for disk rebalancing	No	Yes	No
asm_preferred_read_failure_groups		String	preferred read failure groups	No	Yes	No
audit_file_dest	C:\APP\ADMIN	String	Directory in which auditing files are to reside	Yes	Yes	No
audit_sys_operations	FALSE	Boolean	enable sys auditing	No	No	No
audit_trail	DB	String	enable system auditing	Yes	No	No
background_core_dump	partial	String	Core Size for Background Processes	No	No	No
background_dump_dest	c:\app\m\diag\	String	Detached process dump directory	No	Yes	No
backup_tape_io_slaves	FALSE	Boolean	BACKUP Tape I/O slaves	No	Yes	No
bitmap_merge_area_size	1048576	Integer	maximum memory allow for BITMAP MERGE	No	No	No
blank_trimming	FALSE	Boolean	blank trimming semantics parameter	No	No	No
buffer_pool_keep		String	Number of database blocks/latches in keep bu	No	No	No
buffer_pool_recycle		String	Number of database blocks/latches in recycle	No	No	No
cell_offload_compaction	ADAPTIVE	String	Cell packet compaction strategy	No	Yes	No
cell_offload_decryption	TRUE	Boolean	enable SQL processing offload of encrypted da	No	Yes	No
cell_offload_parameters		String	Additional cell offload parameters	No	Yes	No
cell_offload_plan_display	AUTO	String	Cell offload explain plan display	No	Yes	No
cell_offload_processing	TRUE	Boolean	enable SQL processing offload to cells	No	Yes	No
cell_partition_large_extents	DEFAULT	String	Enables large extent allocation for partitioned t	No	Yes	No
circuits		Integer	max number of circuits	No	Yes	No
client_result_cache_lag	3000	Big integer	client result cache maximum lag in millisecond	No	No	No
client_result_cache_size	0	Big integer	client result cache max size in bytes	No	No	No
cluster_database	FALSE	Boolean	if TRUE startup in cluster database mode	No	No	Yes

**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** , **Table** open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Database** Choose a database to display.

**Report** Choose a report to display.

### Audited Objects

Each row is a different audited object on the selected database. Column values refer to the object.

**Success** This data is obtained from the vendor. See vendor documentation for details.

**Failure** This data is obtained from the vendor. See vendor documentation for details.

**Schema** This data is obtained from the vendor. See vendor documentation for details.

<b>Object Name</b>	This data is obtained from the vendor. See vendor documentation for details.
<b>Statement</b>	This data is obtained from the vendor. See vendor documentation for details.
<b>Object Type</b>	This data is obtained from the vendor. See vendor documentation for details.

**Audited Privileges**

Each row is a different audited privilege on the selected database. Column values refer to the privilege.

<b>Privilege</b>	The name of the privilege.
<b>User</b>	The name of the user associated with the audit.
<b>Proxy</b>	The name of the proxy. This data is obtained from the vendor. See vendor documentation for details.
<b>Success</b>	This data is obtained from the vendor. See vendor documentation for details.
<b>Failure</b>	This data is obtained from the vendor. See vendor documentation for details.

**Database Initialization**

Each row is a different parameter on the selected database. Column values refer to the parameter.

<b>Parameter</b>	The name of the parameter.
<b>Value</b>	The configuration setting of the parameter.
<b>Type</b>	The configuration setting type. Values are: <ul style="list-style-type: none"> <li>• <b>Integer</b></li> <li>• <b>Boolean</b></li> <li>• <b>String</b></li> <li>• <b>Big Integer</b></li> </ul>
<b>Description</b>	Text describing the parameter.
<b>Modified</b>	Indicates whether the parameter was modified. Values are <b>Yes/No</b> . This data is obtained from the vendor. See vendor documentation for details.
<b>Dynamic</b>	Indicates whether the parameter is dynamic. Values are <b>Yes/No</b> . This data is obtained from the vendor. See vendor documentation for details.
<b>Basic</b>	Indicates whether the parameter is basic. Values are <b>Yes/No</b> . This data is obtained from the vendor. See vendor documentation for details.

**Feature Usage**

Each row is a different feature on the selected database. Column values refer to the feature.

<b>Feature Name</b>	The name of the feature.
---------------------	--------------------------

<b>Currently Used</b>	Indicates whether the feature is currently being used. Values are <b>True/False</b> .
<b>#Detected Uses</b>	The total number of times the feature has been used since the database started.
<b>Total Samples</b>	This data is obtained from the vendor. See vendor documentation for details.
<b>First Use</b>	The date and time the feature was first used.
<b>Last Use</b>	The date and time the feature was last used.
<b>Version</b>	This data is obtained from the vendor. See vendor documentation for details.

**Invalid Objects**

Each row is a different invalid object on the selected database. Column values refer to the invalid object.

<b>Owner</b>	The object owner.
<b>Object Type</b>	The type of object.
<b>Object Name</b>	The name of the object.

---

## Oracle Database - HTML

The Oracle Database HTML displays provide extensive visibility into the health and performance of Oracle databases. The HTML version features an ["Oracle Database Overview Display- HTML"](#) (pictured below), and several Views which can be found under **Components** tab > **Databases** > **Oracle Database**:

- ["Oracle Database Overview Display- HTML"](#): This display presents a health snap-shot of the Oracle Database system.
- ["Oracle Databases View - HTML"](#): These displays present metrics about Oracle databases.
- ["Oracle Instances View - HTML"](#): These displays present metrics about instances on Oracle databases.Si

### Oracle Database Overview Display- HTML

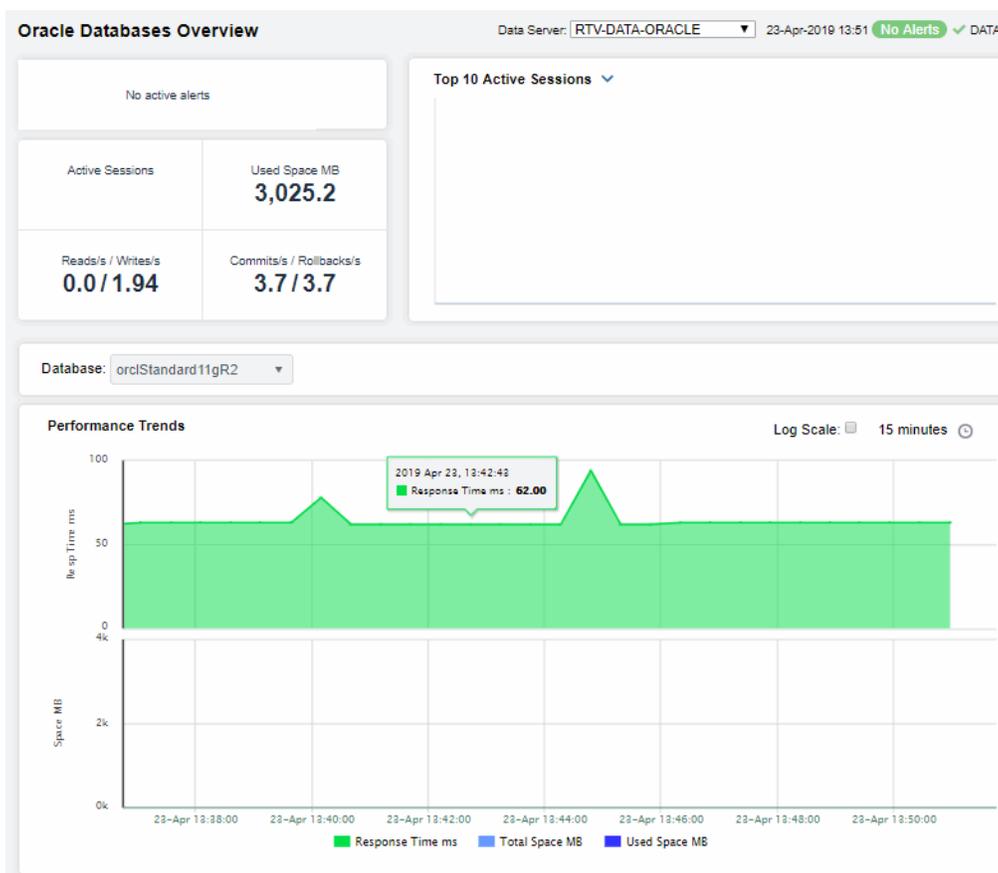
The **Oracle Database Overview** is the top-level display for the Oracle Database Solution Package, which provides a good starting point for immediately getting the status of all your Oracle Databases on your Data Server.

You can select the RTView DataServer for which you want to see data and easily view the current data for that DataServer including:

- The total number of **Active Alerts**, including the total number of critical and warning alerts.
- The number of **Active Sessions** across all databases.
- The total amount of **Used Space** across all databases.
- The number of **Reads / Writes** per second across all databases.
- The number of **Commits / Rollbacks** per second across all databases.
- A bar graphs showing the **Top 10 Active Sessions**.

You can hover over each region in the upper half of the Overview to see more detail. You can also drill-down to see even more detail in the "[Oracle Databases Table - HTML](#)" by clicking on each respective region in the Overview.

The bottom half of the display allows you to select a database for the performance trend graph to trace **Response Time**, **Total Space** and **Used Space**. You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.



## Oracle Databases View - HTML

These displays present performance metrics for Oracle databases. Displays are:

- ["Oracle Databases Table - HTML"](#): A list of all Oracle databases with detailed utilization metrics.
- ["Databases Heatmap - HTML"](#): A heatmap shows alert status of all Oracle databases.
- ["Databases Summary - HTML"](#): Detailed performance metrics for one Oracle database.

### Oracle Databases Table - HTML

Investigate and compare detailed utilization metrics for all Oracle databases. This display contains all metrics available for Oracle databases, including **Space Used**, **Free Space**, **Total Space** and the number of **Instances** on each.

Each row in the table contains data for a particular Oracle database. Click a column header to sort column data in ascending or descending order. Double-click on a table row to drill-down to the ["Databases Summary - HTML"](#) display and view metrics for that particular Oracle database. Toggle between the commonly accessed Heatmap and Summary displays by clicking the drop down list on the display title.

Right-click on a table cell to **Export to Excel** or **Copy Cell Value**.

The **Alert Level** column indicates the most critical alert on the database, where:

- Red indicates that one or more alerts exceeded their ALARM LEVEL threshold in the table row.
- Yellow indicates that one or more alerts exceeded their WARNING LEVEL threshold in the table row.
- Green indicates that no alerts have reached their alert thresholds.

Oracle Databases Table 23-Apr-2019 16:54 No Alerts ✓ DATA

Databases: 1

Database	Connected	Expired	Alert Level	Alert Count	Instances	Resp Time ms	Space Used %	Used Space	Free Space
ora1Standard11gR2	✔		✔		1	83.0	85.08	3,021.25	542.8

## Databases Heatmap - HTML

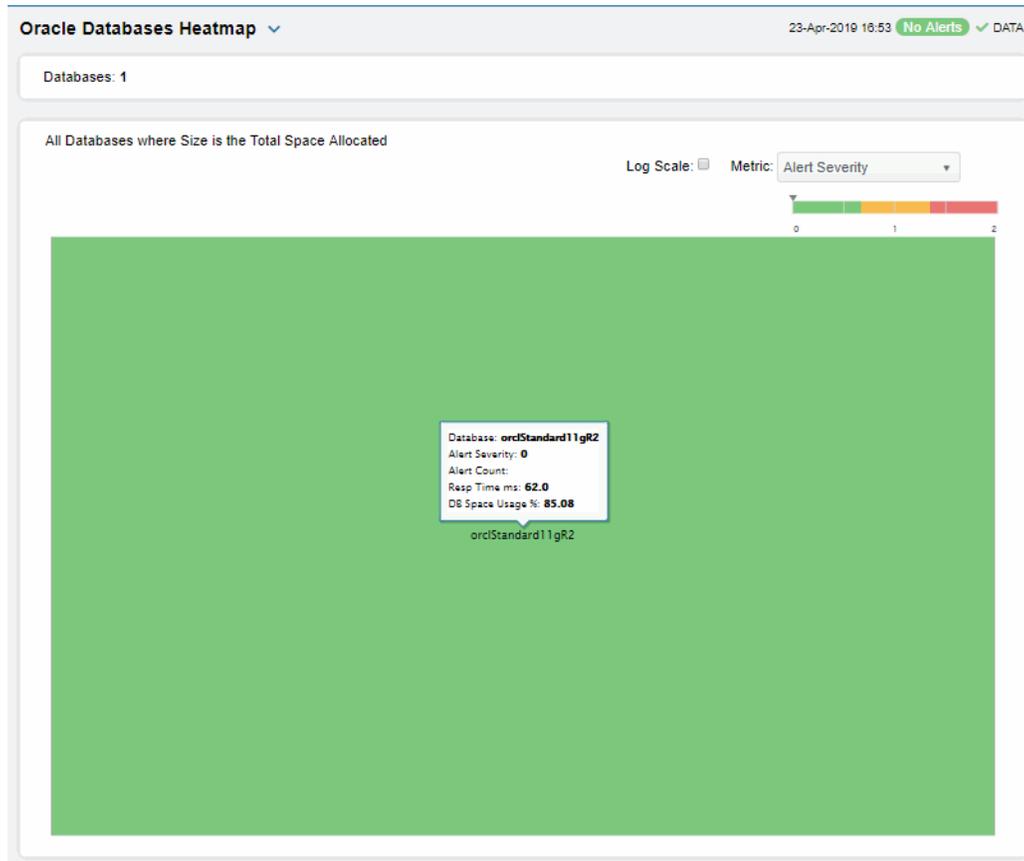
View status and alerts for all Oracle databases. Use this display to quickly identify a database with performance or utilization issues.

Choose a **Metric** from the drop-down menu: **Alert Severity**, **Alert Count**, **Response Time** or **DB Space Usage %**.

Each heatmap rectangle represents a different database. The rectangle color indicates the most critical alert state for the selected metric. Click rectangle to view details in the ["Databases Summary - HTML"](#) display and view metrics for a particular database.

Toggle between the commonly accessed Table and Heatmap displays by clicking the drop down list on the display title. Mouse-over rectangles to view more details about database performance and status.

You can view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.



### Metric:

Choose the type of metric to show in the heatmap. Each metric has its own gradient bar that maps current relative values to colors:

#### Alert Severity

The maximum level of alerts in the heatmap rectangle. Values range from **0** - **2**, as indicated in the color gradient bar, where **2** is the highest Alert Severity.

Red indicates that one or more metrics have reached their alarm threshold. Metrics that have exceeded their specified ALARM LEVEL threshold have an Alert Severity value of **2**.

Yellow indicates that one or more metrics have reached their alarm threshold. Metrics that have exceeded their specified WARNING LEVEL threshold have an Alert Severity value of **1**.

Green indicates that no metrics have reached their alert thresholds. Metrics that have not exceeded their specified thresholds have an Alert Severity value of **0**.

#### Alert Count

The total number of critical and warning alerts in the heatmap rectangle. The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.

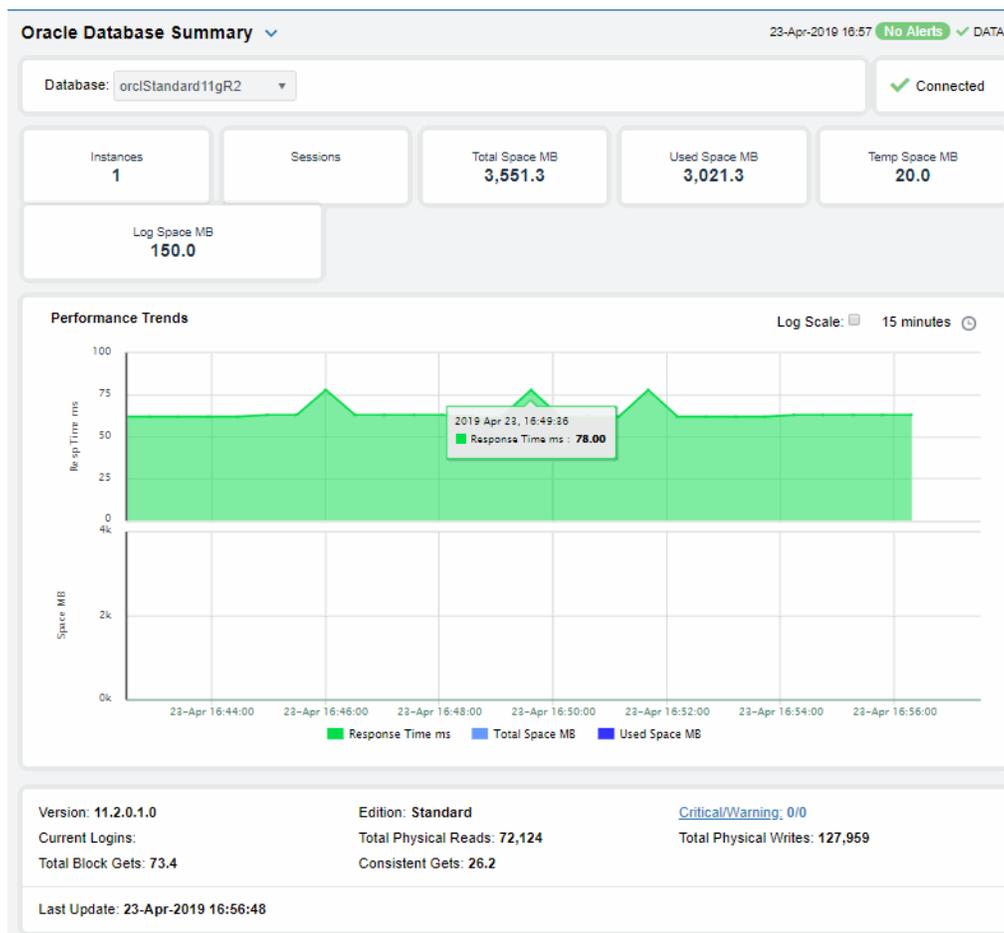
- Response Time** The amount of time, in milliseconds, since the last execution of the database. The color gradient  bar numerical values range from **0** to the maximum amount of time in the heatmap. The middle value in the gradient bar indicates the average amount.
- Database Space Usage** The amount of space used, in megabytes, in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum amount of space used in the heatmap. The middle value in the gradient bar indicates the average amount.

## Databases Summary - HTML

Track utilization and performance metrics for a specific database. Choose a **Database** from the drop-down menu. Mouse-over the utilization information boxes at the top of the display to see more details. Clicking on them takes you to the "[Oracle Databases Table - HTML](#)" display, where you can compare metrics with other Oracle databases.

The trend graph traces **Response Time**, **Total Space** and **Used Space**. Mouse-over to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.

Clicking the **Critical/Warning** link at the bottom of the display opens the Alerts Table by Component display.



## Oracle Instances View - HTML

These displays present performance metrics for Oracle Database instances. Displays are:

- ["Oracle Instances Table - HTML"](#): A list of all Oracle Database instances with detailed utilization metrics.
- ["Instances Heatmap - HTML"](#): A heatmap shows alert status of Oracle Database instances.
- ["Instance Summary - HTML"](#): Detailed performance metrics for one Oracle Database instance.

### Oracle Instances Table - HTML

Investigate and compare detailed utilization and performance metrics as well as configurations for all Oracle Database instances. This display contains all information available for Oracle Database instances such as **Instance ID**, **Instance Name**, **Expired**, **Host Name**, **Instance Role**, **Database Status**, **Version**, **SQL Hit Ratio** and **DD Hit Ratio**.

Choose one or **All** databases from the **Database** drop-down menu. Each row in the table contains data for a particular Oracle Database instance. Click a column header to sort column data in ascending or descending order. Double-click on a table row to drill-down to the ["Instance Summary - HTML"](#) display and view metrics for that particular Oracle Database instance. Toggle between the commonly accessed Heatmap and Summary displays by clicking the drop down list on the display title.

Right-click on a table cell to **Export to Excel** or **Copy Cell Value**.

The **Alert Level** column indicates the most critical alert on the instance, where:

- Red indicates that one or more alerts exceeded their ALARM LEVEL threshold in the table row.
- Yellow indicates that one or more alerts exceeded their WARNING LEVEL threshold in the table row.
- Green indicates that no alerts have reached their alert thresholds.

Oracle Instances Table 24-Apr-2019 09:19 No Alerts ✓ DATA

Database: - All - More Columns

Instances: 1

Database	Instance Id	Instance Name	Expired	Alert Level	Alert Count	Instance Role	DB Status	Version	Active Sessions	Ina
orclStandard11gR2	0	orcl		✓		PRIMARY_INSTAN	ACTIVE	11.2.0.1.0	0.0	

## Instances Heatmap - HTML

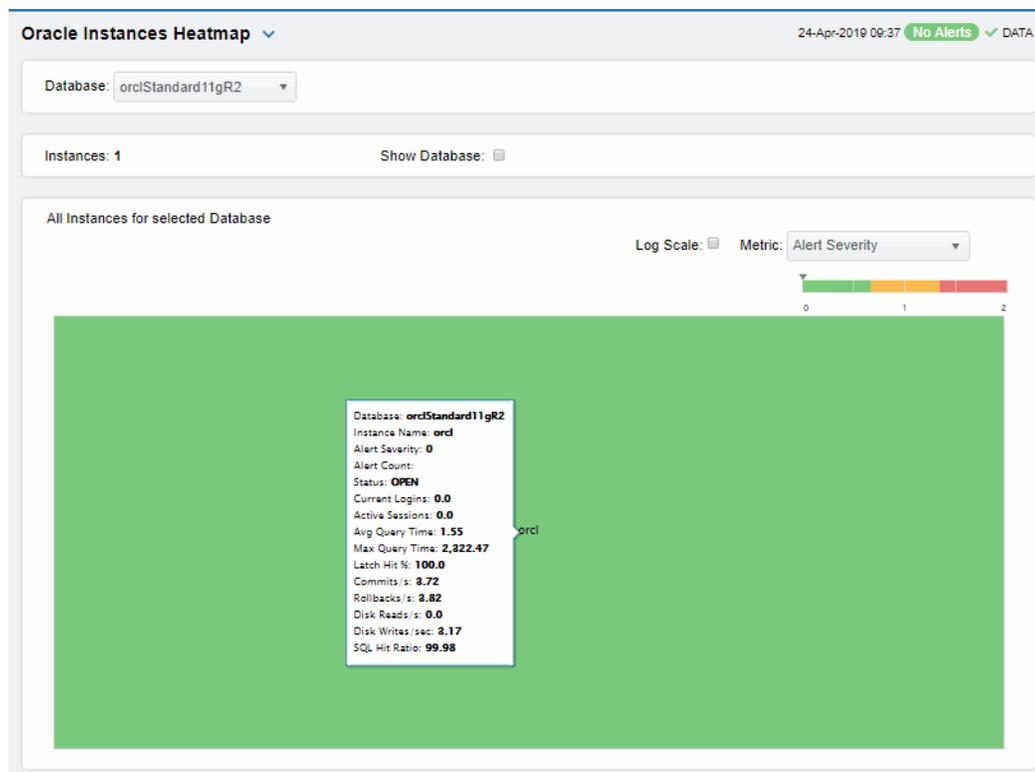
View status and alerts for all instances on one or all Oracle databases. Use this display to quickly identify an instance with performance or utilization issues.

Choose a **Metric** from the drop-down menu: **Alert Severity**, **Alert Count**, **Response Time** or **DB Space Usage %**, **Current Logins**, **Active Sessions**, **Avg Query Time**, **Max Query Time**, **Latch Hit Ratio**, **Data Dict Hit Ratio**, **SQL Hit Ratio**, **Commits/sec**, **Rollbacks/sec**, **Disk Reads/sec** or **Disk Writes/sec**.

Each heatmap rectangle represents a different instance. The rectangle color indicates the most critical alert state for the selected metric. Click rectangle to view details in the ["Instance Summary - HTML"](#) display and view metrics for a particular instance.

Toggle between the commonly accessed Table and Summary displays by clicking the drop down list on the display title. Mouse-over rectangles to view more details about database instance performance and status.

You can view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.



**Database** Choose a database to display.

**Instance Count** The number of instances in the display.

**Metric:**

Choose the type of metric to show in the heatmap. Each rectangle is an instance. Each metric has its own gradient bar that maps current relative values to colors:

<b>Alert Severity</b>	<p>The maximum level of alerts in the heatmap rectangle. Values range from <b>0</b> - <b>2</b>, as indicated in the color gradient  bar, where <b>2</b> is the highest Alert Severity.</p> <ul style="list-style-type: none"> <li> Red indicates that one or more metrics have reached their alarm threshold. Metrics that have exceeded their specified ALARM LEVEL threshold have an Alert Severity value of <b>2</b>.</li> <li> Yellow indicates that one or more metrics have reached their alarm threshold. Metrics that have exceeded their specified WARNING LEVEL threshold have an Alert Severity value of <b>1</b>.</li> <li> Green indicates that no metrics have reached their alert thresholds. Metrics that have not exceeded their specified thresholds have an Alert Severity value of <b>0</b>.</li> </ul>
<b>Current Logins</b>	<p>The number of users logged on. The color gradient  bar values range from <b>0</b> to the maximum count in the heatmap. The middle value in the gradient bar indicates the average count.</p>
<b>Active Sessions</b>	<p>The number of active sessions. The color gradient  bar values range from <b>0</b> to the maximum count in the heatmap. The middle value in the gradient bar indicates the average count.</p>
<b>Avg Query Time</b>	<p>The average amount of time, in seconds, to perform a query. The color gradient  bar values range from <b>0</b> to the maximum average in the heatmap. The middle value in the gradient bar indicates the average amount.</p>
<b>Max Query Time</b>	<p>The maximum amount of time, in seconds, to perform a query. The color gradient  bar values range from <b>0</b> to the maximum amount in the heatmap. The middle value in the gradient bar indicates the average maximum amount.</p>
<b>Latch Hit Ratio</b>	<p>The ratio of the number of latch misses to the number of latch gets. The color gradient  bar values range from the lowest count to the maximum count in the heatmap. The middle value in the gradient bar indicates the average.</p>
<b>Data Dict Hit Ratio</b>	<p>The ratio of logical reads to physical disk reads. The color gradient  bar values range from the lowest count to the maximum value in the heatmap. The middle value in the gradient bar indicates the average.</p>
<b>SQL Hit Ratio</b>	<p>The ratio of logical reads to physical disk reads. The color gradient  bar values range from the lowest count to the maximum value in the heatmap. The middle value in the gradient bar indicates the average.</p>
<b>Commits/sec</b>	<p>The number of commits per second. The color gradient  bar values range from <b>0</b> to the maximum count in the heatmap. The middle value in the gradient bar indicates the average.</p>
<b>Rollbacks/sec</b>	<p>The number of rollbacks per second. The color gradient  bar values range from <b>0</b> to the maximum count in the heatmap. The middle value in the gradient bar indicates the average.</p>
<b>Disk Reads/sec</b>	<p>The number of disk reads per second. The color gradient  bar values range from <b>0</b> to the maximum count in the heatmap. The middle value in the gradient bar indicates the average.</p>
<b>Disk Writes/sec</b>	<p>The number of disk writes per second. The color gradient  bar values range from <b>0</b> to the maximum count in the heatmap. The middle value in the gradient bar indicates the average.</p>

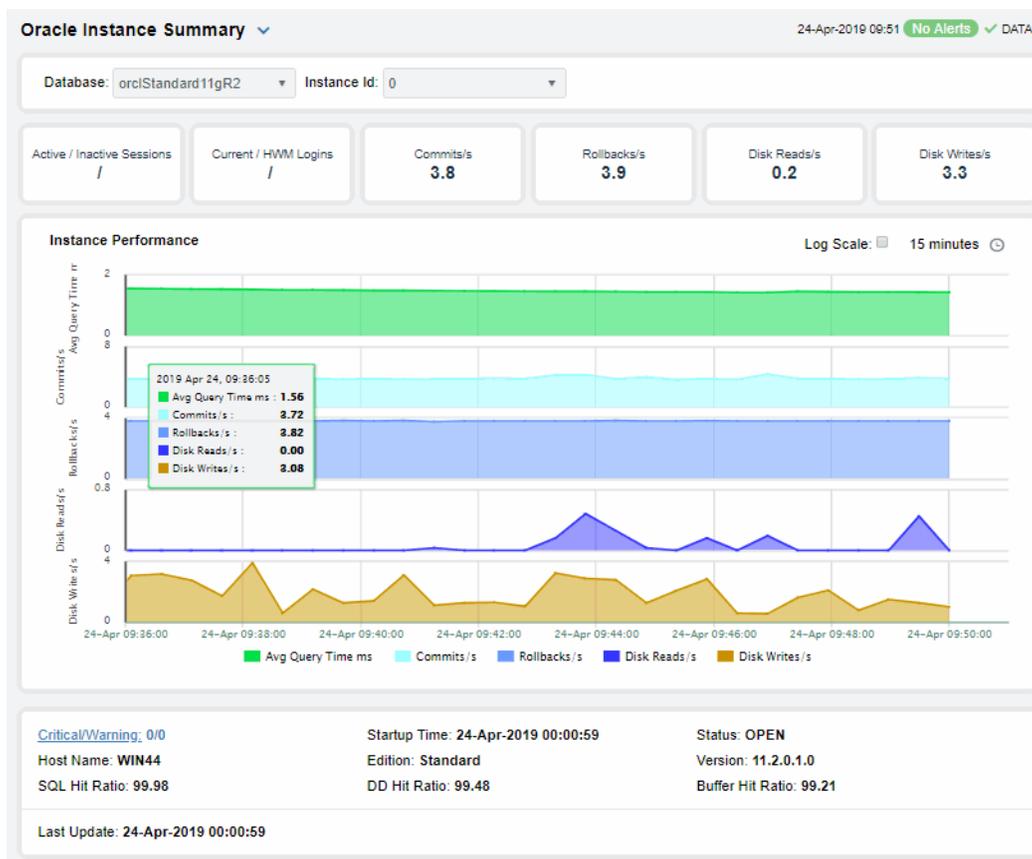
## Instance Summary - HTML

Track utilization and performance metrics for a specific instance. Choose a **Database** and an **Instance** from the drop-down menus. Mouse-over the utilization information boxes at the top of the display to see more details. Clicking on them takes you to the ["Oracle Instances Table - HTML"](#) display, where you can compare metrics with other Oracle Database instances.

The trend graph traces **Avg Query Time** in milliseconds and rates for **Commits**, **Rollbacks**, **Disk Reads** and **Disk Writes**.

Mouse-over to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.

Clicking the **Critical/Warning** link at the bottom of the display opens the Alerts Table by Component display.



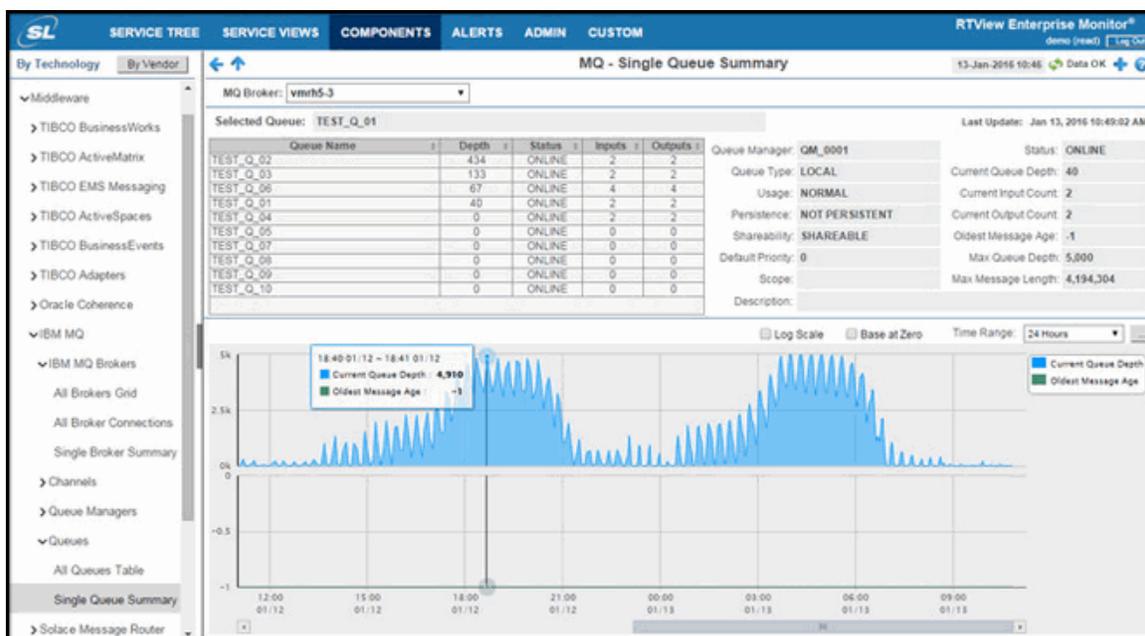
## Connector for Oracle Enterprise Manager

The Connector for Oracle Enterprise Manager (OEM) allows RTView Enterprise to connect to existing deployments of OEM and collect performance data for databases and hosts (physical servers) that have been designated as OEM targets.

When paired with the "Oracle Database" and "RTView Host Agent" solution packages, these performance metrics are then stored in the RTView Enterprise caches and available for summary views detailing the health of your OEM managed hosts and databases, including drill down views, correlation with services and other technologies, historical analysis, capacity planning and alert management.

## Oracle WebLogic

Use the Solution Package for Oracle® WebLogic to monitor the health of your Oracle WebLogic servers and applications as well as your JMS destinations, bridges and connections.



**Note:** This document assumes familiarity with Oracle WebLogic. For additional details, refer to vendor documentation.

The following Solution Package for Oracle® WebLogic Views can be found under **Components** tab > **Application/Web Servers**> **Oracle WebLogic**:

- ["WebLogic Servers View"](#)
- ["Single WebLogic Server View"](#)
- ["Application Views View"](#)
- ["JMS Servers View"](#)
- ["JMS Destinations View"](#)
- ["JMS Bridges View"](#)
- ["JMS Connections View"](#)

## WebLogic Servers View

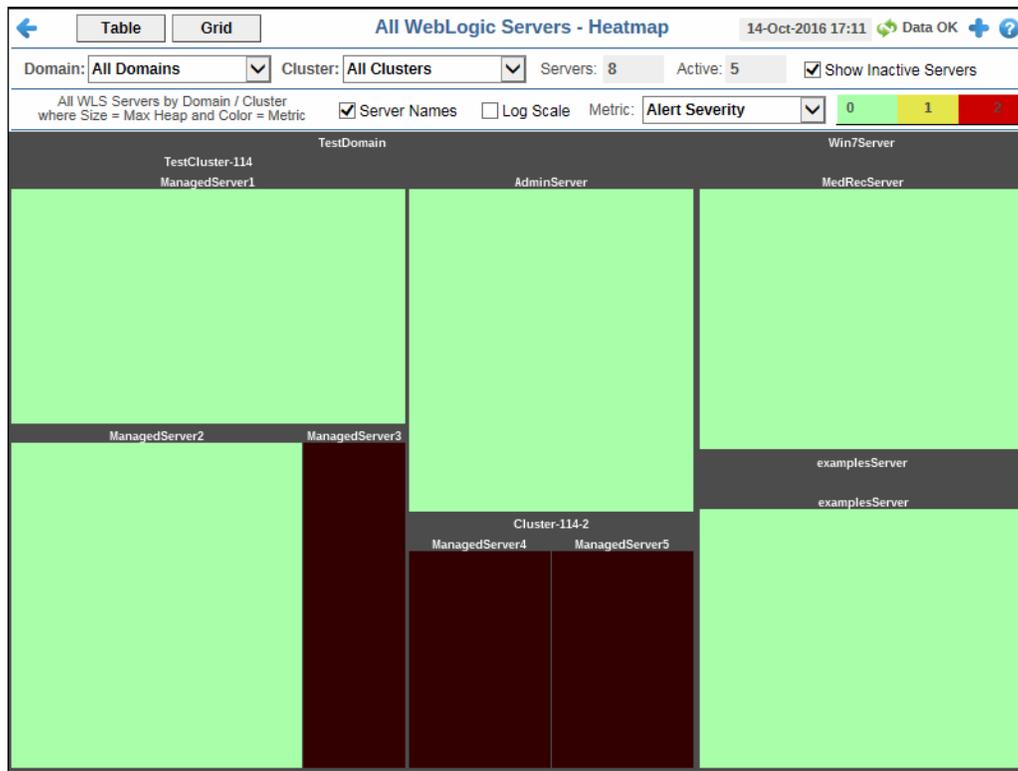
These displays present performance metrics and alert statuses for all Oracle WebLogic Servers and Clusters. The following displays are available:

- ["All Servers Heatmap"](#): This display shows status and alerts for all Oracle® WebLogic servers in a heatmap.
- ["All Servers Table"](#): This displays shows all available utilization metrics for all Oracle® WebLogic servers in a tabular format.
- ["All Servers Grid"](#): This display enables you to track utilization and performance metrics and trend data for all WebLogic servers on a particular domain.
- ["All Clusters Table"](#): This table enables you track utilization and performance metrics for all clusters on a particular domain, or on all domains

### All Servers Heatmap

View status and alerts of all Oracle® WebLogic servers. Use the **Metric** drop-down menu to view the **Alert Severity, Alert Count, Jvm CPU %, Host CPU %, Jvm Memory %, Open Sockets, Thread Total Count, and Hogging Threads**.

The heatmap is organized by host, each rectangle representing a server. The rectangle color indicates the most critical alert state. Click on a node to drill-down to the ["All Servers Grid"](#) display and view metrics for a particular connection. You can toggle between the commonly accessed **Table** and **Heatmap** displays by clicking the icon in the upper left-hand corner. Mouse-over rectangles to view more details about host performance and status.



**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Fields and Data**

This display includes:

- Domain** Select the domain (or **All Domains**) from the drop down list for which you want to view data.
- Cluster** Select the cluster (or **All Clusters**) from the drop down list for which you want to view data.
- Servers** The total number of active, inactive, and standby servers.
- Active** The number of active servers listed in the display.
- Show Inactive Servers** Select this check box to display inactive servers in the heatmap.
- Server Names** Select this check box to display the names of the servers in the heatmap

**Log Scale**

This option enables visualization on a logarithmic scale, and should be used when the range in your data is very broad. For example, if you have data that ranges from the tens to the thousands, then data in the range of tens will be neglected visually if you do not check this option. This option makes data on both extreme ranges visible by using the logarithmic of the values rather than the actual values.

**Metric**

Select the metric driving the heatmap display. The default is Alert Severity. Each **Metric** has a color gradient bar that maps values to colors. The heatmap organizes the instances by host, where each rectangle represents an instance. Mouse-over any rectangle to display the current values of the metrics for the instance. Click on a rectangle to drill-down to the associated "All Servers Grid" display for a detailed view of metrics for that particular instance.

**Alert Severity**

The maximum alert level in the item (index) associated with the rectangle. Values range from **0** to **2**, as indicated in the color gradient bar

 , where **2** is the greatest **Alert Severity**.

**2** -- Metrics that have exceeded their specified **ALARMLEVEL** threshold and have an Alert Severity value of **2** are shown in red. For a given rectangle, this indicates that one or more metrics have exceeded their alarm threshold.

**1** -- Metrics that have exceeded their specified **WARNINGLEVEL** threshold and have an Alert Severity value of **1** are shown in yellow. For a given rectangle, this indicates that one or more metrics have exceeded their warning threshold.

**0** -- Metrics that have not exceeded either specified threshold have an Alert Severity value of **0** and are shown in green. For a given rectangle, this indicates that no metrics have exceeded a specified alert threshold.

**Alert Count**

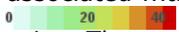
The total number of alarm and warning alerts in a given item (index) associated with the rectangle.

The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

**Jvm CPU %**

The percentage of JVM CPU currently being used in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the alert threshold of **WlsServerCpuHigh**. The middle value in the gradient bar indicates the middle value of the range.

**Host CPU %**

The percentage of Host CPU currently being used in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the alert threshold of **WlsServerHostCpuHigh**. The middle value in the gradient bar indicates the middle value of the range.

<b>JVM Memory %</b>	The percentage of JVM Memory currently being used in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the alert threshold of <b>WisMemoryUsageHigh</b> . The middle value in the gradient bar indicates the middle value of the range.
<b>Open Sockets</b>	The total number of open sockets currently being used in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the alert threshold of <b>WisOpenSocketsHigh</b> . The middle value in the gradient bar indicates the middle value of the range.
<b>Thread Total Count</b>	The total number of threads in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the alert threshold of <b>WisThreadsTotalHigh</b> . The middle value in the gradient bar indicates the middle value of the range.
<b>Hogging Threads</b>	The total number of hogging threads currently being used in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the alert threshold of <b>WisHoggingThreadsHigh</b> . The middle value in the gradient bar indicates the middle value of the range.

## All Servers Table

This display provides utilization metrics for all WebLogic Servers for a particular domain in a tabular format. Each row in this table includes heap, processing, thread, and version metrics (among others) for a particular server. Click a column header to sort column data in numerical or alphabetical order. Click on a table row to drill-down to the ["Server Summary"](#) display and view metrics for that particular server. You can click on one of the buttons in the upper left-hand corner to toggle between the commonly accessed **Grid** and **Heatmap** displays.

Domain	Cluster	Server	State	Expired	Alerts	Port	JVM Proc Load %	Heap Free %	H
examplesServer		examplesServer	RUNNING	<input type="checkbox"/>		80	7.21	61.0	5
TestDomain		AdminServer	RUNNING	<input type="checkbox"/>		6001	1.60	5.0	5
TestDomain	Cluster-114-2	ManagedServer4	SHUTDO...	<input type="checkbox"/>		0	0	0.0	5
TestDomain	Cluster-114-2	ManagedServer5	SHUTDO...	<input type="checkbox"/>		0	0	0.0	5
TestDomain	TestCluster-114	ManagedServer1	RUNNING	<input type="checkbox"/>		6003	1.26	31.0	5
TestDomain	TestCluster-114	ManagedServer2	RUNNING	<input type="checkbox"/>		6004	0.00	47.0	5
TestDomain	TestCluster-114	ManagedServer3	SHUTDO...	<input type="checkbox"/>		0	0	0.0	5
Win7Server		MedRec Server	RUNNING	<input type="checkbox"/>		80	2.99	76.0	5

#### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- The number of items currently in the display.

- Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

**Note:** Fields with an asterisk (\*) at the end of the field definition contain data that is provided by the WebLogic Server MBean interface. Refer to WebLogic documentation for more information regarding these fields.

### Fields and Data

This display includes:

- Domain** Select the domain for which you want to view data, or select **All Domains** to view data for all domains.
- Cluster** Select the cluster on the domain for which you want to view data, or select **All Clusters** to view data for all domains.
- Servers** The total number of servers on the cluster.
- Active** The total number of active servers on the cluster.

**Show Inactive Servers**

Select this check box to display inactive servers in the table.

**All WebLogic Servers Table**

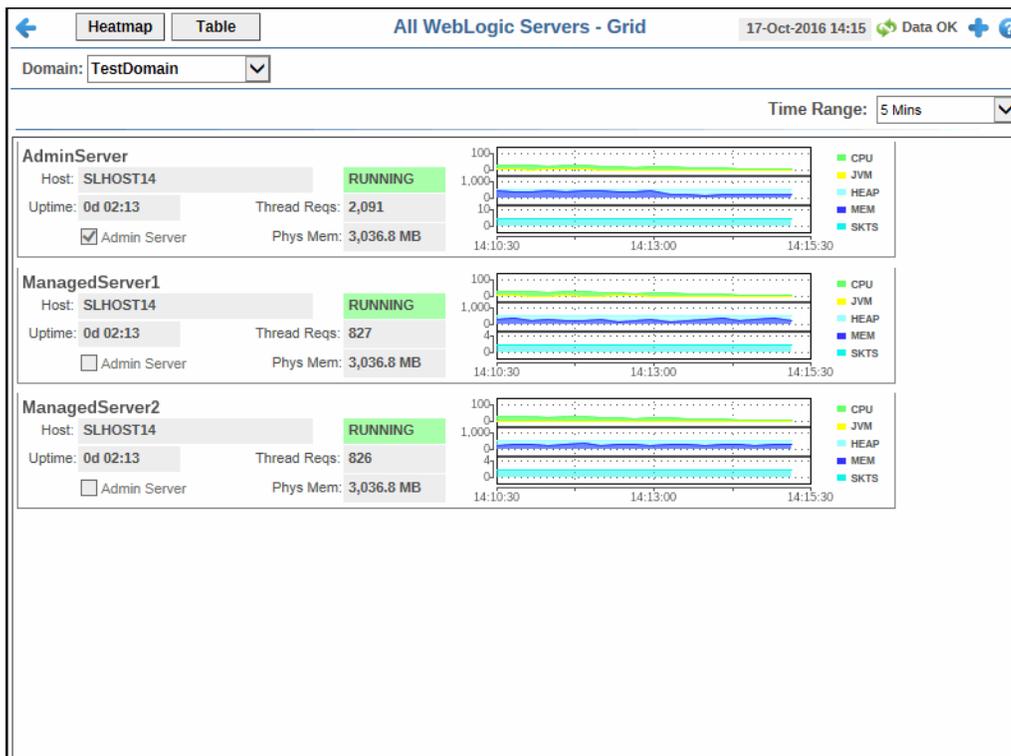
This table shows information for the selected domain/cluster(s) combination. Click on a table row to drill-down to the ["Server Summary"](#) display and view metrics for that particular server.

<b>Domain</b>	The name of the domain.
<b>Cluster</b>	The name of the cluster.
<b>Server</b>	The name of the server.
<b>State</b>	The current state of the server.
<b>Expired</b>	This check box becomes automatically checked when the data displayed in the row has exceeded the specified cache expiration time (set by default at 45 seconds) and is no longer current. Once the cache has been refreshed and is displaying current data, the check box will return to being unchecked. This check box will remain unchecked as long as the cache has been refreshed within the specified cache expiration time and the data is current.
<b>Alerts</b>	The current alert level.  -- One or more alerts have exceeded their specified <b>ALARMLEVEL</b> threshold.  -- One or more alerts have exceeded their specified <b>WARNINGLEVEL</b> threshold.  -- No alerts have exceeded an alert threshold.
<b>Port</b>	The port on which this server is listening for SSL connections.*
<b>JVM Proc Load %</b>	A snapshot of the load that the virtual machine is placing on all processors on the host computer.*
<b>Heap Free %</b>	The percentage of free heap memory on the server.*
<b>Heap Max (bytes)</b>	The maximum amount of heap, in bytes, available for use.*
<b>Used Heap (bytes)</b>	The total amount of heap used, in bytes.*
<b>Heap Current (bytes)</b>	The current size of the JVM heap, in bytes, being used.*
<b>Open Sockets</b>	The current number of sockets registered for socket muxing on this server.*
<b>Hogging Threads</b>	The number of hogging threads on the server.*
<b>Execute Threads</b>	The current number of execute threads.*
<b>Idle Threads</b>	The current number of idle threads.
<b>Restarts Count</b>	The total number of restarts for this server since the cluster was last started.*
<b>All Procs Avg Load %</b>	The average load percentage for all processors on the host computer.*
<b>Shutting Down</b>	When checked, denotes that the server is currently shutting down.*

<b>Restart Required</b>	When checked, denotes that the server needs to be restarted in order to activate configuration changes.*
<b>Uptime</b>	The length of time (in milliseconds) that the server has been up and running.*
<b>Startup Time</b>	The length of time (in milliseconds) that it took for the server to start up.*
<b>WebLogic Version</b>	The current version of WebLogic running on the server.*
<b>JVM Type</b>	The type of JVM currently being used on the server.*
<b>Java Version</b>	The current version of Java running on the server.*
<b>JavaVendor</b>	The name of the vendor of the Java version running on the server.*
<b>OS Name</b>	The name of the operating system running on the server.*
<b>Time Stamp</b>	The date and time this row of data was last updated.

### All Servers Grid

Track utilization and performance metrics and trend data for all WebLogic servers on a particular domain.



**Title Bar** (possible features are):

-   Open the previous and upper display.
-  Open an instance of this display in a new window.
-  Open the online help page for this display.
- ,  open commonly accessed displays.
- The number of items currently in the display.

 **Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

 Open the **Alert Views - RTView Alerts Table** display.

---

**Note:** Fields with an asterisk (\*) at the end of the field definition contain data that is provided by the WebLogic Server MBean interface. Refer to WebLogic documentation for more information regarding these fields.

---

### Fields and Data

This display includes:

<b>Domain</b>	Select the domain for which you want to view data, or select <b>All Domains</b> to view data for all domains.
<b>Time Range</b>	Select a time range from this drop down menu to define the data displayed in the trend graph for a selected period of time. You can select from as little as the past <b>2 Minutes</b> to the <b>Last 7 Days</b> , or you can display <b>All Data</b> .
<b>Server Grid</b>	Displays data and a trend graph for each server in your domain. The trend graph metrics are: <ul style="list-style-type: none"> <li><b>CPU</b> -- Traces the amount of CPU being used by the server.</li> <li><b>JVM</b> -- Traces the JVM processing load that the virtual machine is placing on all processors on the host computer.</li> <li><b>HEAP</b> -- Traces the total amount of heap used, in bytes.</li> <li><b>MEM</b> -- Traces the amount of memory being used.</li> <li><b>SKTS</b> -- Traces the current number of sockets registered for socket muxing on this server.</li> </ul>
<b>(Server Name)</b>	Displays the name of the Server
<b>Uptime</b>	The amount of time since the server was last started, shown in days, hours, and minutes (for example, 1d 23:43).
<b>Admin Server</b>	Indicates whether the server is an Administration Server.*
<b>(Status)</b>	Displays the status of the server.*
<b>Thread Reqs</b>	The current number of thread requests.*
<b>Phys Mem</b>	Displays the available physical memory (in megabytes) for the server.

## All Clusters Table

Track utilization and performance metrics for all clusters on a particular domain, or on all domains.

Domain	Cluster	Server Count	Running Count	Percent Not Running	Time Stamp
examplesServer		1	1	0.0	17-Oct-2016 14:27:52
TestDomain	TestCluster-114	3	2	33.3	17-Oct-2016 14:27:52
TestDomain	Cluster-114-2	2	0	100.0	17-Oct-2016 14:27:52
TestDomain		1	1	0.0	17-Oct-2016 14:27:52
Win7Server		1	1	0.0	17-Oct-2016 14:27:52

### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- The number of items currently in the display.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

Open the **Alert Views - RTView Alerts Table** display.

**Note:** Fields with an asterisk (\*) at the end of the field definition contain data that is provided by the WebLogic Server MBean interface. Refer to WebLogic documentation for more information regarding these fields.

### Fields and Data

This display includes:

- Domain** Select the domain containing the clusters for which you want to view data, or select **All Domains** to view data for clusters in all domains.

**Cluster Count** The current number of clusters listed in the table.

#### Clusters Table

Lists the clusters in the currently selected domain, or lists all clusters in all domains.

<b>Domain</b>	The name of the domain
<b>Cluster</b>	The name of the cluster.*
<b>Server Count</b>	The total number of servers on the cluster.*
<b>Running Count</b>	The total number of servers running on the cluster.*
<b>Percent Not Running</b>	The percentage of servers not running on the cluster.*
<b>Time Stamp</b>	The date and time this row of data was last updated.

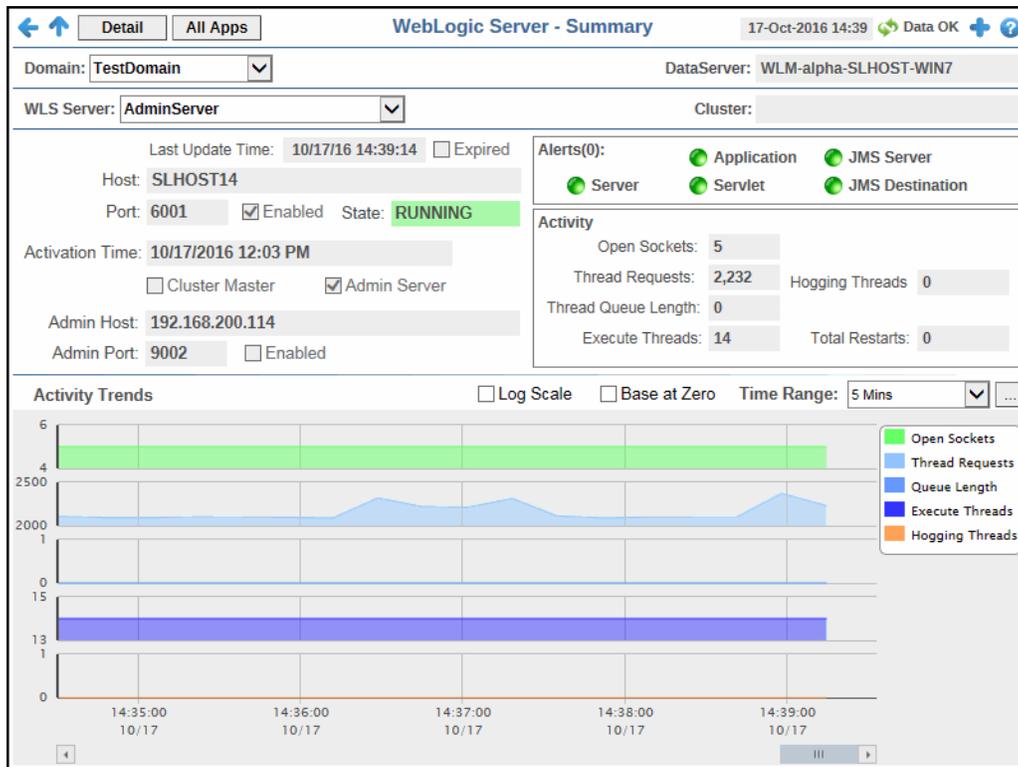
## Single WebLogic Server View

These displays present detailed performance metrics and alert statuses for a particular WebLogic server.

- ["Server Summary"](#): Track utilization, performance, and trend data for a particular WebLogic server
- ["WLS JVM Summary"](#): Displays the JVM details for a particular WebLogic server on a specific domain.
- ["WLS Server Detail"](#): Displays server runtime data, threadpool runtime data, JRockit runtime data, and server version information for a specific WebLogic server
- ["WLS JDBC Summary"](#): Displays JDBC module utilization, performance, and trend data for a specific WebLogic server.
- ["WLS ThreadPool Summary"](#): Displays threadpool utilization, performance, and trend data for a specific WebLogic server.
- ["Work Manager"](#): Displays server runtime data for all work managers on a specific WebLogic Server.
- ["Persistent Stores"](#): Displays available utilization and performance data for all configurations on a specific domain.

## Server Summary

Track utilization, performance, and trend data for a particular WebLogic server.



**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- 6,047** The number of items currently in the display.
- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

**Note:** Fields with an asterisk (\*) at the end of the field definition contain data that is provided by the WebLogic Server MBean interface. Refer to WebLogic documentation for more information regarding these fields.

### Fields and Data

This display includes:

- Domain** Select the domain for which you want to view data.

<b>Data Server</b>	The name of the data server.
<b>WLS Server</b>	Select the WebLogic server for which you want to see data.
<b>Cluster</b>	The name of the cluster.

**(Server Information)**

<b>Last Update Time</b>	The date and time the data in the display was last updated.
<b>Expired</b>	This check box becomes automatically checked when the data has exceeded the specified cache expiration time (set by default at 45 seconds) and is no longer current. Once the cache has been refreshed and is displaying current data, the check box will return to being unchecked. This check box will remain unchecked as long as the cache has been refreshed within the specified cache expiration time and the data is current.
<b>Host</b>	The name of the host.*
<b>Port/Enabled</b>	The name of the port. The port is enabled when the associated <b>Enabled</b> check box displays as checked.*
<b>State</b>	The current state of the WebLogic server.*
<b>Activation Time</b>	The date and time in which the server was started.*
<b>Cluster Master</b>	When selected, denotes that the server is a cluster master.*
<b>Admin Server</b>	Indicates whether the server is an Administration Server.*
<b>Admin Host</b>	The IP address of the administration server's host.*
<b>Admin Port/Enabled</b>	The name of the administration server's port. The port is enabled when the associated <b>Enabled</b> check box displays as checked.*

**Alerts (#)** -- Displays the total number of alerts and the current status of the associated **Application, JMS Server, Server, Servlet, and JMS Destination.**

-  -- One or more alerts have exceeded their specified **ALARMLEVEL** threshold.
-  -- One or more alerts have exceeded their specified **WARNINGLEVEL** threshold.
-  -- No alerts have exceeded an alert threshold.

**Activity**

<b>Open Sockets</b>	The number of current open sockets for the server.
<b>Thread Requests</b>	The current number of thread requests.*
<b>Hogging Threads</b>	The current number of hogging threads.*
<b>Thread Queue Length</b>	The current thread queue length.*
<b>Execute Threads</b>	The current number of execute threads.*
<b>Total Restarts</b>	The total number of times the server has restarted since the last update time.*

**Activity Trends**

Displays data and a trend graph for the following:

- Open Sockets**-- Traces the number of open sockets of the server.
- Thread Requests**-- Traces the number of thread requests on the server.
- Queue Length**-- Traces the queue length on the server.
- Execute Threads**-- Traces the number of execute threads on the server.
- Hogging Threads**-- Traces the number of hogging threads on the server.

**Log Scale**

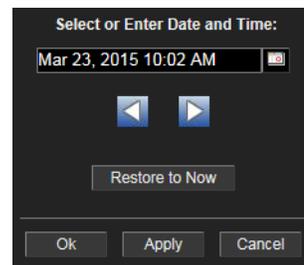
This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.

**Base at Zero**

When this option is checked, zero is set as the Y axis minimum for all graph traces.

**Time Range**

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



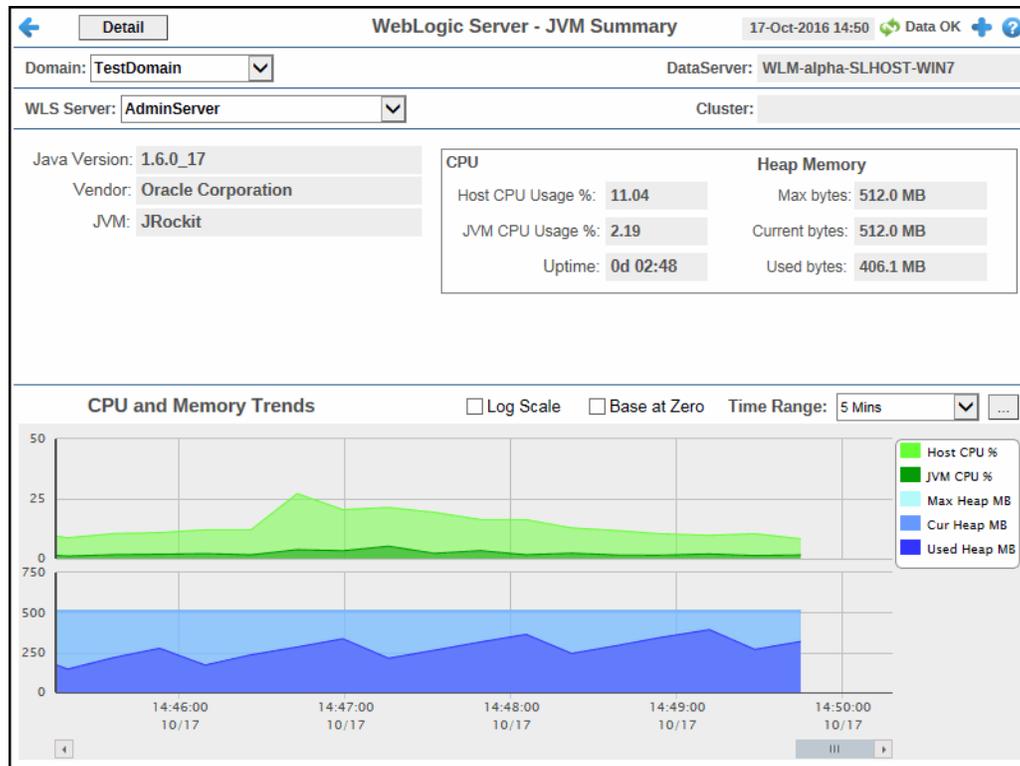
By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## WLS JVM Summary

This display allows you to view the JVM details for a particular WebLogic server on a specific domain.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- The number of items currently in the display.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

Open the Alert Views - RTView Alerts Table display.

**Note:** Fields with an asterisk (\*) at the end of the field definition contain data that is provided by the WebLogic Server MBean interface. Refer to WebLogic documentation for more information regarding these fields.

### Fields and Data

This display includes:

**Domain** Select the domain for which you want to view data.

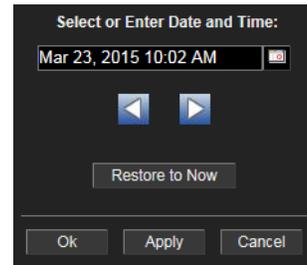
<b>Data Server</b>	The name of the data server.
<b>WLS Server</b>	Select the WebLogic server for which you want to see data.
<b>Cluster</b>	The name of the cluster.
<b>(JVM Information)</b>	
<b>Java Version</b>	The current version of Java running on the server.*
<b>Vendor</b>	The name of the vendor of the current version of Java running on the server.*
<b>JVM</b>	The type of JVM currently being used on the server.*
<b>CPU</b>	
<b>Host CPU Usage %</b>	The current CPU usage percentage on the host.*
<b>JVM CPU Usage %</b>	The current JVM CPU usage percentage.*
<b>Uptime</b>	The amount of time since the server was last started, shown in days, hours, and minutes (for example, 1d 23:43).*
<b>Heap Memory</b>	
<b>Max Bytes</b>	The maximum amount of available heap memory, in megabytes.*
<b>Current Bytes</b>	The current size of the JVM heap, in megabytes.*
<b>Used Bytes</b>	The amount of heap memory used, in megabytes.*
<b>Alerts (#)</b> -- Displays the total number of alerts and the current status of the associated <b>Application, JMS Server, Server, Servlet,</b> and <b>JMS Destination.</b>	
● -- One or more alerts have exceeded their specified <b>ALARMLEVEL</b> threshold.	
● -- One or more alerts have exceeded their specified <b>WARNINGLEVEL</b> threshold.	
● -- No alerts have exceeded an alert threshold.	
<b>CPU and Memory Trends</b>	Displays data and a trend graph for the following: <ul style="list-style-type: none"> <li><b>Host CPU %</b> -- Traces the percentage of the host CPU being used.</li> <li><b>JVM CPU %</b> -- Traces the percentage of the JVM CPU being used.</li> <li><b>Max Heap MB</b> -- Traces the maximum amount of heap memory available (in megabytes).</li> <li><b>Current Heap MB</b> -- Traces the current size of the JVM heap, in megabytes.</li> <li><b>Used Heap MB</b> -- Traces the total amount of heap used, in bytes.*</li> </ul>
<b>Log Scale</b>	This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.

**Base at Zero**

When this option is checked, zero is set as the Y axis minimum for all graph traces.

**Time Range**

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## WLS Server Detail

View server runtime, threadpool runtime, JRockit runtime, and server version information for a specific WebLogic server.

The screenshot shows the 'WLS Server - Detail Tables' interface. At the top, there is a 'Summary' tab and a title bar with the date '17-Oct-2016 14:59' and a 'Data OK' icon. Below the title bar, there are dropdown menus for 'Domain: TestDomain', 'WLS Server: AdminServer', and 'DataServer: WLM-alpha-SLHOST-WIN7'. The main content area is divided into four sections, each with a table and a scroll bar:

- Server Runtime:** A table with columns: Connection, Location, Cluster, ActivationTime, AdminServer, AdminServerHost, AdminServerListenPort, and Ad. The data row shows: TestDomain, AdminServer, 1476731039316, [checked], 192.168.200.114, 6001.
- ThreadPool Runtime:** A table with columns: Connection, CompletedRequestCount, DeltaCompletedRequestCount, ExecuteThreadIdleCount, ExecuteThreadTotalCount, and Hod. The data row shows: TestDomain, 1,353,303, 2,098, 8, 14.
- JRockit Runtime:** A table with columns: Connection, AllProcessorsAverageLoad, Concurrent, FreeHeap, FreePhysicalMemory, GCHandlesCompaction, and Dynamic GC. The data row shows: TestDomain, 6.934, [unchecked], 215,658,408, 929,218,560, [checked], Dynamic GC.
- Server Version Info:** A text area containing: WebLogic Server 10.3.3.0 Fri Apr 9 00:05:28 PDT 2010 1321401.

### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- The number of items currently in the display.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

Open the **Alert Views - RTView Alerts Table** display.

**Note:** Row columns in this table with an asterisk (\*) at the end of the column definition contain data that is provided by the WebLogic Server MBean interface. Refer to WebLogic documentation for more information regarding these columns.

### Fields and Data

This display includes:

**Domain** Select the domain for which you want to view data.

<b>Data Server</b>	The name of the data server.
<b>WLS Server</b>	Select the WebLogic server for which you want to see data.
<b>Cluster</b>	The name of the cluster.
<b>Server Runtime Table</b>	
<b>Connection</b>	The name of the connection.
<b>Location</b>	The name of the WLS Server located on the specified connection.
<b>Cluster</b>	The name of the cluster.
<b>Activation Time</b>	The time when the server was started.*
<b>Admin Server</b>	Indicates whether the server is an Administration Server.*
<b>Admin Server Host</b>	The address on which the Administration Server is listening for connections.*
<b>Admin Server Listen Port</b>	The port on which the Administration Server is listening for connections.*
<b>Admin Server Listen Port Secure</b>	Indicates whether the port that the server uses for administrative traffic is configured to use a secure protocol.*
<b>Administration Port</b>	The port on which this server is listening for administrative requests.*
<b>Administration Port Enabled</b>	Indicates whether the administration port is enabled on the server.*
<b>Administration URL</b>	The URL that the server and its clients use for administrative connections.*
<b>Cluster Master</b>	When checked, denotes that the cluster is a cluster master.*
<b>Current Directory</b>	The absolute path of the directory from which the server was started.*
<b>Current Machine</b>	The machine on which the server is running.*
<b>Default URL</b>	The URL that clients use to connect to this server's default network channel.*
<b>Health State</b>	The health state of the server as reported by the server's self-health monitoring.*
<b>Listen Address</b>	The address on which this server is listening for connections through the default network channel.*
<b>Listen Port</b>	The port on which this server is listening for connections.*
<b>Listen Port Enabled</b>	Indicates whether the default listen port is enabled on the server.*
<b>Name</b>	The name of the Java Virtual Machine.*
<b>Open Sockets Current Count</b>	The current number of sockets registered for socket muxing on this server.*

<b>Oracle Home</b>	The Oracle home directory path.*
<b>Parent</b>	The name of the parent of the Java Virtual Machine.*
<b>Pending Restart System Resources</b>	The number of system resources that have not been restarted since the last restart of the WLS Server.
<b>Restart Required</b>	Indicates whether the server must be restarted in order to activate configuration changes.*
<b>Restarts Total Count</b>	The total number of restarts for this server since the cluster was last started.*
<b>SSL Listen Address</b>	The address on which this server is listening for SSL connections.*
<b>SSL Listen Port</b>	The port on which this server is listening for SSL connections.*
<b>SSL Listen Port Enabled</b>	Indicates whether the default SSL listen port is enabled on the server.*
<b>Server Classpath</b>	The class path for this server including domain/lib contents that are automatically picked up and appended to the classpath.*
<b>Server Startup Time</b>	The startup time of the server.*
<b>Shutting Down</b>	Indicates whether the server is shutting down.
<b>Stable State</b>	The current state of the server as an integer.*
<b>State Prev</b>	The state of the server prior to its current state.*
<b>State Val</b>	The current state of the server as an integer.*
<b>WebLogic Home</b>	The WebLogic Home directory path.*
<b>Expired</b>	This check box becomes automatically checked when the data displayed in the row has exceeded the specified cache expiration time (set by default at 45 seconds) and is no longer current. Once the cache has been refreshed and is displaying current data, the check box will return to being unchecked. This check box will remain unchecked as long as the cache has been refreshed within the specified cache expiration time and the data is current.
<b>State</b>	The current life cycle state of the server.*

**ThreadPool Runtime Table**

<b>Connection</b>	The name of the connection.
<b>Completed Request Count</b>	The number of completed requests in the priority queue.*
<b>Delta Completed Request Count</b>	The increase in the amount of completed requests (from the previous polling period to the current polling period).

<b>Execute Thread Idle Count</b>	The number of idle threads in the pool. This count does not include standby threads and stuck threads. The count indicates threads that are ready to pick up new work when it arrives.*
<b>Execute Thread Total Count</b>	The total number of threads in the pool.*
<b>Hogging Thread Count</b>	The threads that are currently being hogged by a request. These threads will either be declared as stuck after the configured timeout or will be returned to the pool. The self-tuning mechanism will backfill if necessary.*
<b>Delta Hogging Thread Count</b>	The increase in the amount of hogging threads (from the previous polling period to the current polling period).
<b>Min Threads Constraints Completed</b>	The number of requests with minimum threads constraint picked up out of order for execution immediately since their minimum threads requirement was not met. This does not include the case where threads are idle during schedule.*
<b>Min Threads Constraints Pending</b>	The number of requests that should be executed now to satisfy the minimum threads requirement.*
<b>Pending User Request Count</b>	The number of pending user requests in the priority queue. The priority queue contains requests from internal subsystems and users. This is just the count of all user requests.*
<b>Queue Length</b>	The number of pending requests, which consist of the total number of internal system requests and user requests, in the priority queue.*
<b>Shared Capacity For Work Managers</b>	The maximum amount of requests that can be accepted in the priority queue.*
<b>Standby Thread Count</b>	The number of threads in the standby pool. Surplus threads that are not needed to handle the present work load are designated as standby and added to the standby pool. These threads are activated when more threads are needed.*
<b>Suspended</b>	Indicates if the RequestManager is suspended. A suspended manager will not dequeue work and dispatch threads until it is resumed.*
<b>Throughput</b>	The mean number of requests completed per second.*
<b>Name</b>	The name of the Java Virtual Machine.
<b>Parent</b>	The name of the parent of the Java Virtual Machine.

#### JRockit Runtime Table

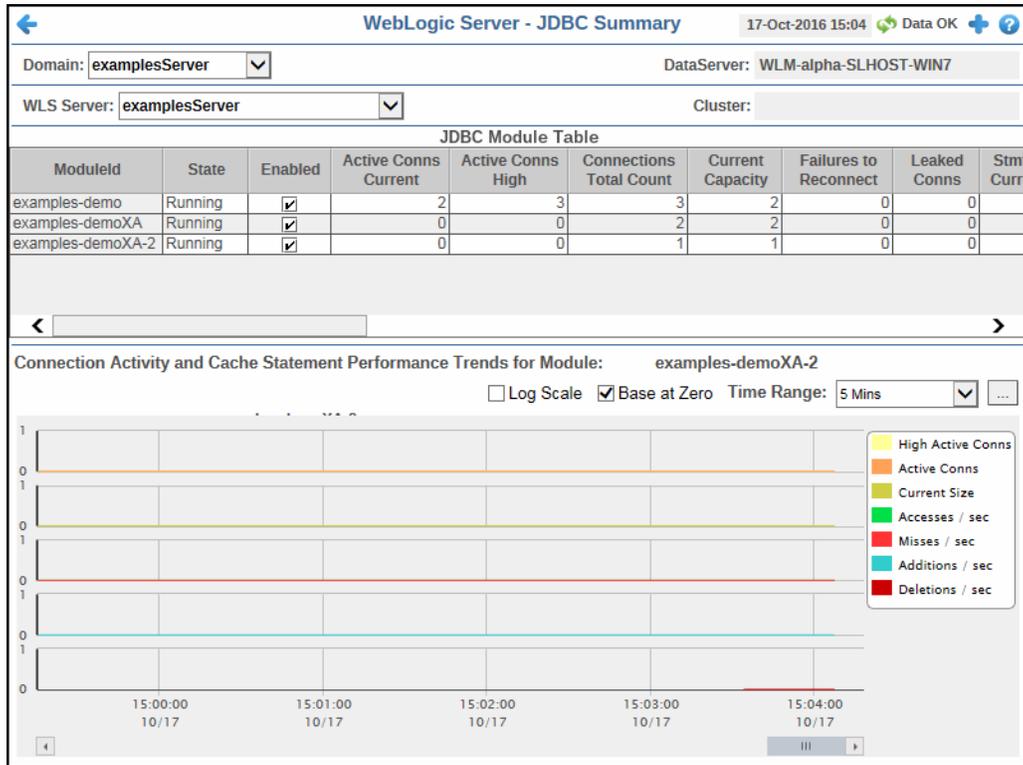
<b>Connection</b>	The name of the connection.
<b>All Processors Average Load</b>	The average load of all processors in the host computer.*
<b>Concurrent</b>	Indicates whether the virtual machine's garbage collector runs in a separate Java thread concurrently with other Java threads.*

<b>Free Heap</b>	The amount, in bytes, of Java heap memory that is currently free in the virtual machine.*
<b>Free Physical Memory</b>	The amount, in bytes, of physical memory that is currently free on the host computer.*
<b>GC Handles Compaction</b>	Indicates whether the virtual machine's garbage collector compacts the Java heap.*
<b>GcAlgorithm</b>	The type of garbage collector (GC) that the virtual machine is using.*
<b>Generational</b>	Indicates whether the virtual machine's garbage collector uses a nursery space. A nursery is the area of the Java heap that the virtual machine allocates to most objects.*
<b>Heap Free Current</b>	The current amount of memory, in bytes, that is available in the JVM heap.*
<b>Heap Free Percent</b>	Percentage of the maximum memory that is free.*
<b>Heap Size Current</b>	The current size, in bytes, of the JVM heap.*
<b>Heap Size Max</b>	The maximum free memory configured for this JVM.*
<b>Incremental</b>	Indicates whether the virtual machine's garbage collector collects (increments) garbage as it scans the memory space and dumps the garbage at the end of its cycle. With a non-incremental garbage collector, garbage is dumped as soon as it is encountered.*
<b>JVM Description</b>	The description of the Java Virtual Machine.*
<b>Java VM Vendor</b>	The vendor of the Java Virtual Machine that the server is running.*
<b>Java Vendor</b>	The vendor of Java that the server is running.*
<b>Java Version</b>	The Java version of the Java Virtual Machine.*
<b>JVM Processor Load</b>	A snapshot of the load that the virtual machine is placing on all processors in the host computer. If the host contains multiple processors, the value represents a snapshot of the average load.*
<b>Jvm Type</b>	The Java Virtual Machine type.*
<b>Last GC End</b>	The time at which the last garbage collection run ended.*
<b>Last GC Start</b>	The time at which the last garbage collection run started.*
<b>Name</b>	The name of the Java Virtual Machine.*
<b>Number Of Daemon Threads</b>	The number of daemon Java threads currently running in the Virtual Machine across all processors.*

<b>Number Of Processors</b>	The number of processors on the virtual machine's host computer. If this is not a Symmetric Multi-Processor (SMP) system, the value will be 1.*
<b>OS Name</b>	The name of the operating system on which the JVM is running.*
<b>OS Version</b>	The version of the operating system on which the JVM is running.*
<b>Parallel</b>	Indicates whether the virtual machine's garbage collector is able to run in parallel on multiple processors if multiple processors are available.*
<b>Parent</b>	The name of the immediate parent.*
<b>Total Garbage Collection Count</b>	The number of garbage collection runs that have occurred since the virtual machine was started.*
<b>Total Garbage Collection Time</b>	The number of milliseconds that the virtual machine has spent on all garbage collection runs since the virtual machine was started.*
<b>Total Heap</b>	The amount, in bytes, of memory currently allocated to the virtual machine's Java heap.*
<b>Total Number Of Threads</b>	The number of Java threads (daemon and non-daemon) that are currently running in the virtual machine across all processors.*
<b>Total Nursery Size</b>	The amount, in bytes, of memory that is currently allocated to the nursery.*
<b>Total Physical Memory</b>	The amount (in bytes) of physical memory on the host computer.*
<b>Uptime</b>	The amount of time, in milliseconds, that the virtual machine has been running.*
<b>Used Heap</b>	The amount of Java heap memory, in bytes, that is currently being used by the virtual machine.*
<b>Used Physical Memory</b>	The amount of physical memory, in bytes, that is currently being used on the host computer.*
<b>Vendor</b>	The name of the JVM vendor.*
<b>Version</b>	The current version of the Java Virtual Machine.*
<b>Heap Used Current</b>	The current amount of JVM heap memory, in bytes, that is being used.*
<b>Memory Used Percent</b>	The percentage of JVM heap memory that is being used.*
<b>Server Version Info Region</b>	Lists the WebLogic server version number and date that it was installed.

## WLS JDBC Summary

View JDBC module utilization, performance data, and trend data for a specific WebLogic server.



**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.
- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

**Note:** Row columns in this table with an asterisk (\*) at the end of the column definition contain data that is provided by the WebLogic Server MBean interface. Refer to WebLogic documentation for more information regarding these columns.

### Fields and Data

This display includes:

- Domain** Select the domain for which you want to view data.
- Data Server** The name of the data server.

**WLS Server** Select the WebLogic server for which you want to see data.

**Cluster** The name of the cluster.

### JDBC Module Table

<b>Module Id</b>	The name of the data source.
<b>State</b>	The current state of the data source.
<b>Enabled</b>	Indicates whether the data source is enabled or disabled. When checked, the data source is enabled.*
<b>Active Conns Current</b>	The number of connections currently in use by applications.*
<b>Active Conns High</b>	The highest number of active database connections in the instance of the data source since the data source was instantiated.*
<b>Connections Total Count</b>	The cumulative total number of database connections created in this data source since the data source was deployed.*
<b>Current Capacity</b>	The current count of JDBC connections in the connection pool in the data source.*
<b>Failures to Reconnect</b>	The number of times that the data source attempted to refresh a database connection and failed.*
<b>Leaked Conns</b>	The number of leaked connections. *
<b>Stmt Cache Current Size</b>	The number of prepared and callable statements currently cached in the statement cache.*
<b>Stmt Cache Hits</b>	The cumulative, running count of the number of times statements from the cache were used.*
<b>Stmt Cache Missed</b>	The number of times that a statement request could not be satisfied with a statement from the cache.*
<b>Stmt Cache Accesses</b>	The cumulative, running count of the number of times that the statement cache was accessed.*
<b>Stmt Cache Additions</b>	The cumulative, running count of the number of statements added to the statement cache.*
<b>Stmt Cache Deletions</b>	The cumulative, running count of statements discarded from the cache.*
<b>Reserve Request</b>	The cumulative, running count of requests for a connection from this data source.*
<b>Failed Reserve Requests</b>	The cumulative, running count of requests for a connection from this data source that could not be fulfilled.*
<b>Wait Secs High</b>	The cumulative total number of database connections created in this data source since the data source was deployed.*
<b>Waiting Conn</b>	The number of connection requests waiting for a database connection.*
<b>Waiting Conn Fail</b>	The cumulative, running count of requests for a connection from this data source that had to wait before getting a connection and eventually failed to get a connection.*

<b>Waiting Conn High</b>	The highest number of application requests concurrently waiting for a connection from this instance of the data source.*
<b>Waiting Conn Success</b>	The cumulative, running count of requests for a connection from this data source that had to wait before getting a connection and eventually succeeded in getting a connection.*
<b>Waiting Conn Total</b>	The cumulative, running count of requests for a connection from this data source that had to wait before getting a connection, including those that eventually got a connection and those that did not get a connection.*
<b>Connection Delay Time</b>	The average amount of time, in milliseconds, that it takes to create a physical connection to the database.*
<b>Driver Version</b>	The driver class name of the JDBC driver used to create database connections.*
<b>time_stamp</b>	The date and time that the data in the row was last updated.*

**Connection Activity and Cache Statement Performance Trends for Module: (Module Name)**

Shows connection and open cursor data for the connection.

**High Active Conns** -- Traces the highest number of active database connections in the instance of the data source since the data source was instantiated.

**Active Conns** -- Traces the number of connections currently in use by applications.

**Current Size** -- Traces the number of prepared and callable statements currently cached in the statement cache.

**Accesses/sec** -- Traces the cumulative, running count of the number of times that the statement cache was accessed.

**Misses/sec** -- Traces the number of times (per second) that a statement request could not be satisfied with a statement from the cache.

**Additions/sec** -- Traces the cumulative, running count of the number of statements added to the statement cache.

**Deletions/sec** -- Traces the cumulative, running count of statements discarded from the cache.

**Log Scale**

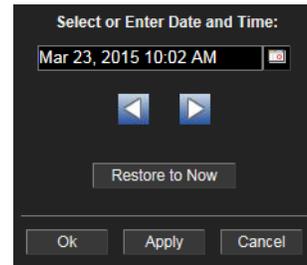
This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.

**Base at Zero**

When this option is checked, zero is set as the Y axis minimum for all graph traces.

**Time Range**

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



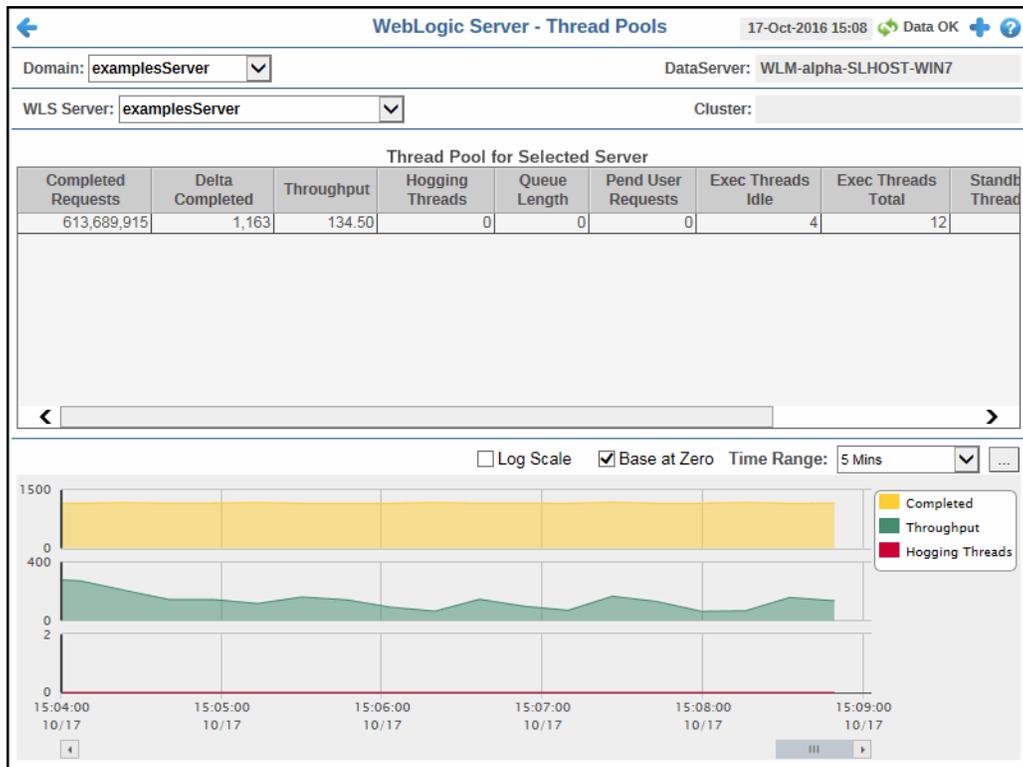
By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## WLS ThreadPool Summary

View threadpool utilization, performance, and trend data for a specific WebLogic server.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- The number of items currently in the display.

- Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

**Note:** Row columns in this table with an asterisk (\*) at the end of the column definition contain data that is provided by the WebLogic Server MBean interface. Refer to WebLogic documentation for more information regarding these columns.

### Fields and Data

This display includes:

- Domain** Select the domain for which you want to view data.
- Data Server** The name of the data server.

**WLS Server** Select the WebLogic server for which you want to see data.

**Cluster** The name of the cluster.

### Thread Pool for Selected Server Table

<b>Completed Requests</b>	The number of completed requests in the priority queue.*
<b>Delta Completed</b>	The increase in the amount of completed requests (from the previous polling period to the current polling period).
<b>Throughput</b>	The mean number of requests completed per second.*
<b>Hogging Threads</b>	The threads that are currently being hogged by a request. These threads will either be declared as stuck after the configured timeout or will be returned to the pool. The self-tuning mechanism will backfill if necessary.*
<b>Queue Length</b>	The number of pending requests, which consist of the total of internal system requests and user requests, in the priority queue.*
<b>Pend User Requests</b>	The number of pending user requests in the priority queue.*
<b>Exec Threads Idle</b>	The number of idle threads in the pool.*
<b>Exec Threads Total</b>	The total number of threads in the pool.*
<b>Standby Threads</b>	The number of threads in the standby pool.*
<b>Suspended</b>	Indicates if the RequestManager is suspended. A suspended manager will not dequeue work and dispatch threads until it is resumed.*
<b>time_stamp</b>	The date and time the data in the row was last updated.

**Trend Graph** Shows connection and open cursor data for the connection.

**Completed** -- Traces the number of completed requests in the priority queue.

**Throughput** -- Traces the mean number of requests completed per second.

**Hogging Threads**-- Traces the number of threads that are currently being hogged by a request.

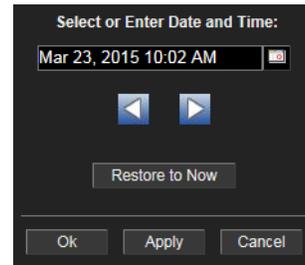
**Log Scale** This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.

**Base at Zero**

When this option is checked, zero is set as the Y axis minimum for all graph traces.

**Time Range**

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Work Manager

View server runtime data for all work managers on a specific WebLogic Server.

The screenshot shows the 'WLS Server - Work Manager' interface. At the top, there is a 'Summary' tab and a title bar with the text 'WLS Server - Work Manager' and a timestamp '17-Oct-2016 15:20'. Below the title bar, there are dropdown menus for 'Domain: examplesServer', 'DataServer: WLM-alpha-SLHOST-WIN7', 'WLS Server: examplesServer', and 'Cluster:'. The main content is a table titled 'Server Runtime' with the following columns: Name, Application, Completed, Delta Compl..., Rate Compl..., Pending, and Stuck. The table lists various work managers and their associated applications, with numerical values in the 'Completed', 'Delta Compl...', 'Rate Compl...', 'Pending', and 'Stuck' columns. For example, 'DataRetirementWorkManager' has 0 in all these columns, while 'JmsAsyncQueue' has 70,624 in the 'Completed' column. The table is scrollable, with arrows visible at the bottom.

**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.
- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

**Note:** Row columns in this table with an asterisk (\*) at the end of the column definition contain data that is provided by the WebLogic Server MBean interface. Refer to WebLogic documentation for more information regarding these columns.

### Fields and Data

This display includes:

- Domain** Select the domain for which you want to view data.
- Data Server** The name of the data server.

**WLS Server** Select the WebLogic server for which you want to see data.

**Cluster** The name of the cluster.

#### Server Runtime Table

<b>Name</b>	The name of the work manager.
<b>Application</b>	The name of the application with which the work manager is associated.*
<b>Completed</b>	The number of requests that have been completed.*
<b>Delta Completed</b>	The increase in the amount of completed requests (from the previous polling period to the current polling period).
<b>Rate Completed</b>	The rate of completed requests (per second).
<b>Pending</b>	The number of requests waiting in the queue.*
<b>Stuck</b>	The number of threads that are "stuck."*
<b>Expired</b>	This check box becomes automatically checked when the data displayed in the row has exceeded the specified cache expiration time (set by default at 45 seconds) and is no longer current. Once the cache has been refreshed and is displaying current data, the check box will return to being unchecked. This check box will remain unchecked as long as the cache has been refreshed within the specified cache expiration time and the data is current.
<b>time_stamp</b>	The date and time the data in the row was last updated.

## Persistent Stores

View available utilization and performance data for all configurations on a specific domain.

← **Persistent Stores - Detail Tables** 18-Nov-2016 14:50 Data OK

Domain:

Persistent Store Runtime						
Name	Location	HealthState	CreateCo...	DeleteCount	ObjectCo...	Physic
WseeFileStore_auto_1	AdminServer	HEALTH_OK	1	1	9	
WseeJaxwsFileStore_auto_1	ManagedServer1	HEALTH_OK	2	2	21	
WseeJaxwsFileStore_auto_2	ManagedServer2	HEALTH_OK	2	2	21	

Persistent Store Connection Runtime				
Name	Location	PersistentStoreRuntime	CreateCo...	DeleteCo...
weblogic.messaging.JMServer-1.body	ManagedServer1	FileStore-1	0	0
weblogic.messaging.WseeJmsServer_auto_1.body	AdminServer	WseeFileStore_auto_1	0	0
weblogic.transaction.tlog	ManagedServer2	_WLS_ManagedServer2	0	0

**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- The number of items currently in the display.
- Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

**Note:** Row columns in this table with an asterisk (\*) at the end of the column definition contain data that is provided by the WebLogic Server MBean interface. Refer to WebLogic documentation for more information regarding these columns.

### Fields and Data

This display includes:

**Domain**            Select the domain for which you want to view data.

#### Persistent Store Runtime Table

<b>Name</b>	The name of the configuration.
<b>Location</b>	The name of the WLS Server located on the specified connection.
<b>Health State</b>	The health state of the store.*
<b>Create Count</b>	The number of create requests issued to the persistent store.*
<b>Delete Count</b>	The number of delete requests issued by this persistent store.*
<b>Object Count</b>	The number of objects contained in the persistent store.*
<b>Physical Write Count</b>	The number of times the persistent store flushes its data to durable storage.*
<b>Read Count</b>	The number of read requests issued to this persistent store.*
<b>Update Count</b>	The number of update requests issued by this persistent store.*
<b>time_stamp</b>	The date and time the data in the row was last updated.

#### **Persistent Store Connection Runtime Table**

<b>Name</b>	The name of the configuration.
<b>Location</b>	The name of the WLS Server located on the specified connection.
<b>Persistent Store Runtime</b>	The name of the persistent store.
<b>Create Count</b>	The number of create requests issued by this connection.*
<b>Delete Count</b>	The number of delete requests issued by this connection.*
<b>Object Count</b>	The number of objects contained in this connection.*
<b>Read Count</b>	The number of read requests issued by this connection.*
<b>Update Count</b>	The number of update requests issued by this connection.*
<b>time_stamp</b>	The date and time the data in the row was last updated.

## Application Views View

These displays present several views of performance metrics for applications on clusters and a particular WebLogic server.

- **"Cluster Apps Table"**: View performance and utilization metrics for all applications on a particular cluster, or for all applications on all clusters.
- **"Cluster App Summary"**: View session information for a particular application in graph, heatmap, and trend formats.
- **"Server Apps Heatmap"**: Shows a heatmap view of the status and alerts of all applications within a specific WebLogic server.
- **"Server Apps Summary"**: Track performance, utilization, and trend data for all applications on a single WebLogic server.
- **"Server Apps Trends"**: View trend data for a single application on a particular WebLogic server.
- **"App Components Heatmap"**: Provides a heatmap view of the status and alerts of all application components contained within each application on a particular WebLogic server.
- **"App Components Summary"**: View performance, utilization, and trend data for all application components on a single WebLogic Server.

### Cluster Apps Table

View performance and utilization metrics for all applications on a particular cluster, or for all applications on all clusters.

WebLogic Cluster Applications - Table								17-Oct-2016 15:30	Data OK	+	?
Domain: All Domains		Cluster: All Clusters		Cluster Count: 5							
Domain	Cluster	Application	Sessions	New	Sessions	Invocations	Exec Time				
TestDomain		bea_wls9_async_response	0	0	0.00	0.00	0.00				
TestDomain		bea_wls_deployment_internal	0	0	0.00	0.00	0.00				
TestDomain		bea_wls_internal	0	0	0.00	0.00	0.00				
TestDomain		bea_wls_management_inter...	0	0	0.00	0.00	0.00				
TestDomain	Cluster-114-2		0	0							
TestDomain	TestCluster-114	bea_wls9_async_response	0	0	0.00	0.00	0.00				
TestDomain	TestCluster-114	bea_wls_cluster_internal	0	0	0.00	0.00	0.00				
TestDomain	TestCluster-114	bea_wls_deployment_internal	0	0	0.00	0.28	0.14				
TestDomain	TestCluster-114	bea_wls_internal	0	0	0.00	0.00	0.00				
TestDomain	TestCluster-114	wlm	0	0	0.00	0.00	0.00				
Win7Server		bea_wls9_async_response	0	0	0.00	0.00	0.00				
Win7Server		bea_wls_internal	0	0	0.00	0.00	0.00				
Win7Server		esphere	0	0	0.00	0.00	0.00				
Win7Server		myocm_rtvdata	0	0	0.00	0.00	0.00				
Win7Server		ocmon	0	0	0.00	0.00	0.00				
Win7Server		ocmon_rtvdata	0	0	0.00	0.00	0.00				
Win7Server		ocmon_rtvquery	0	0	0.00	0.00	0.00				
Win7Server		wlm	0	0	0.00	0.00	0.00				
Win7Server		wlm_rtvdata	0	0	0.00	0.00	0.00				
Win7Server		wlm_rtvquery	0	0	0.00	0.00	0.00				
examplesServer		SamplesSearchWebApp	0	0	0.00	0.00	0.00				
examplesServer		asyncServletEar	0	0	0.00	0.00	0.00				
examplesServer		bea_wls9_async_response	0	0	0.00	0.00	0.00				
examplesServer		bea_wls_internal	0	0	0.00	0.00	0.00				
examplesServer		ejb20BeanMgedEar	0	0	0.00	0.00	0.00				
examplesServer		ejb30	0	0	0.00	0.00	0.00				
examplesServer		esphere	0	0	0.00	0.00	0.00				
examplesServer		examplesWebApp	0	0	0.00	0.00	0.00				
examplesServer		extServletAnnotationsEar	0	0	0.00	0.00	0.00				
examplesServer		jdbcRowSetsEar	0	0	0.00	0.00	0.00				
examplesServer		ispSimpleTanEar	0	0	0.00	0.00	0.00				

**Title Bar** (possible features are):

-   Open the previous and upper display.
-  Open an instance of this display in a new window.
-  Open the online help page for this display.
- ,  open commonly accessed displays.
- The number of items currently in the display.

 **Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

 Open the **Alert Views - RTView Alerts Table** display.

---

**Note:** Fields with an asterisk (\*) at the end of the field definition contain data that is provided by the WebLogic Server MBean interface. Refer to WebLogic documentation for more information regarding these fields.

---

### Fields and Data

This display includes:

<b>Domain</b>	Select the domain for which you want to view data.
<b>Cluster</b>	The name of the cluster.
<b>Cluster Count</b>	Displays the total number of clusters listed in the table.

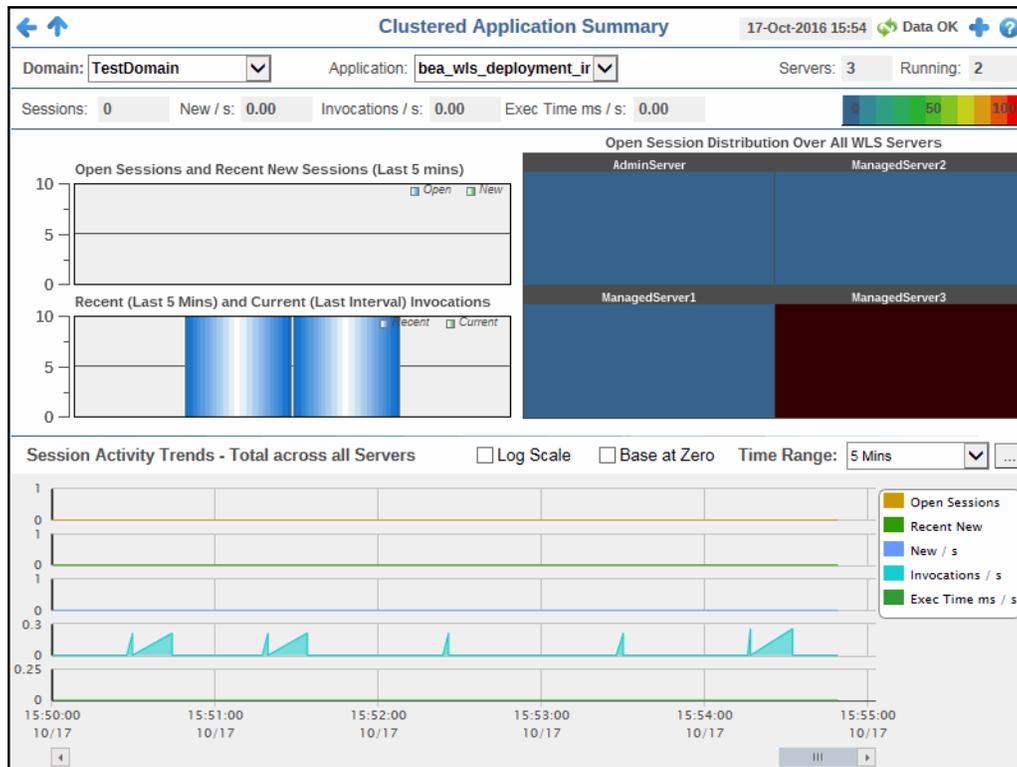
### WebLogic Cluster Applications Table

Click a row to view metrics for a single topic in the "[Cluster App Summary](#)" display.

<b>Domain</b>	The name of the domain.
<b>Cluster</b>	The name of the cluster.
<b>Application</b>	The name of the application.
<b>Sessions Open</b>	The number of open sessions on the application.
<b>New Sessions</b>	The number of new sessions since the last polling update.
<b>Session New/sec</b>	The rate of new sessions (per second).
<b>Invocations/sec</b>	The rate of invocations (per second).
<b>Exec Time ms/sec</b>	The rate of execution time in milliseconds (per second).
<b>Server Count</b>	The total number of existing servers on the application.
<b>Running Count</b>	The total number of running servers on the application.
<b>Percent Not Running</b>	The percentage of servers that are not running.
<b>time_stamp</b>	The date and time this row of data was last updated.

## Cluster App Summary

View session information for a particular application in graph, heatmap, and trend formats.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- 6,047** The number of items currently in the display.

- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

**Note:** Fields with an asterisk (\*) at the end of the field definition contain data that is provided by the WebLogic Server MBean interface. Refer to WebLogic documentation for more information regarding these fields.

### Fields and Data

This display includes:

- Domain** Select the domain containing the application for which you want to view data.

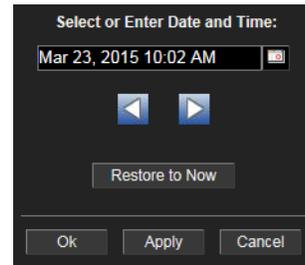
<b>Application</b>	Select the application for which you want to view data.
<b>Servers</b>	Displays the total number of servers on the application.*
<b>Running</b>	Displays the total number of servers running on the application.*
<b>Sessions</b>	Displays the number of open sessions.*
<b>New/s</b>	Displays the rate of new sessions being opened (per second).*
<b>Invocations/s</b>	Displays the rate of invocations (per second).*
<b>Exec Time ms/s</b>	Displays the rate of execution time in milliseconds (per second).*
<b>Open Sessions and Recent New Sessions (Last 5 mins) bar graph</b>	Displays the currently open sessions and the sessions created in the last 5 minutes.
<b>Recent (Last 5 Mins) and Current (Last Interval) Invocation bar graph</b>	Displays the number of recent invocations (last 5 minutes) and the number of current invocations (created since the last polling interval).
<b>Open Session Distribution Over All WLS Servers heat map</b>	Displays the number of open sessions for each WLS server in heatmap form based on the color gradient bar  .
<b>Trend Graph</b>	Shows connection and open cursor data for the connection. <ul style="list-style-type: none"> <li><b>Open Sessions</b> -- Traces the number of open sessions.</li> <li><b>Recent New</b> -- Traces the number of newly created sessions.</li> <li><b>New/s</b> -- Traces the number of sessions created per second.</li> <li><b>Invocations/s</b> -- Traces the number of invocations per second.</li> <li><b>Exec Time ms/s</b> -- Traces the execution time in milliseconds per second.</li> </ul>
<b>Log Scale</b>	This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.

**Base at Zero**

When this option is checked, zero is set as the Y axis minimum for all graph traces.

**Time Range**

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



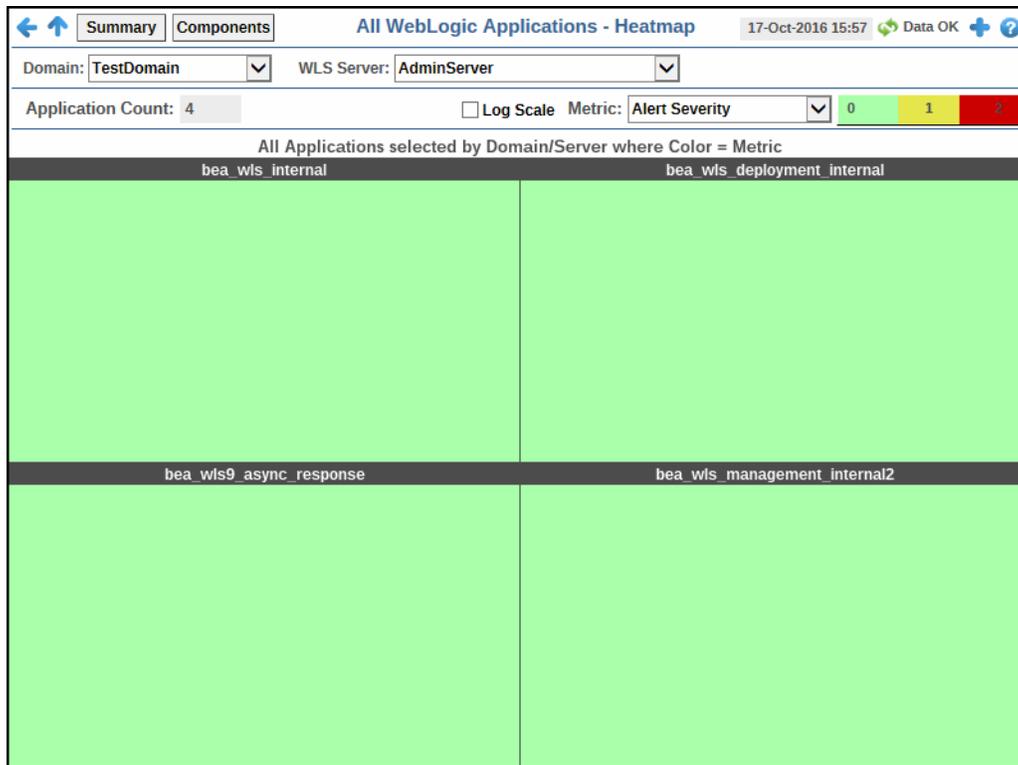
By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows  to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Server Apps Heatmap

This display provides a heatmap view of the status and alerts of all applications within a specific WebLogic server. The heatmap is organized so that each rectangle represents a collection contained within a specific connection. The rectangle color indicates the value of the selected metric in the **Metric** drop down list. You can mouse-over rectangles to view more details about the performance and status of each collection or click on a rectangle to drill-down to the "Server Apps Trends" display and view metrics for that particular collection. You can click the table icon  in this display to navigate to the "Server Apps Summary" display.



### Title Bar (possible features are):

-   Open the previous and upper display.
-  Open an instance of this display in a new window.
-  Open the online help page for this display.
-   open commonly accessed displays.
-  6,047 The number of items currently in the display.

-  Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
-  23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
-  Open the **Alert Views - RTView Alerts Table** display.

**Note:** Fields with an asterisk (\*) at the end of the field definition contain data that is provided by the WebLogic Server MBean interface. Refer to WebLogic documentation for more information regarding these fields.

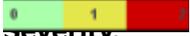
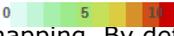
## Fields and Data

This display includes:

<b>Domain</b>	Select the domain containing the WebLogic Server for which you want to view data.
<b>WLS Server</b>	Select the WebLogic server for which you want to view data.
<b>Application Count</b>	The total number of applications on the server.
<b>Log Scale</b>	This option enables visualization on a logarithmic scale, and should be used when the range in your data is very broad. For example, if you have data that ranges from the tens to the thousands, then data in the range of tens will be neglected visually if you do not check this option. This option makes data on both extreme ranges visible by using the logarithmic of the values rather than the actual values.

## Metric

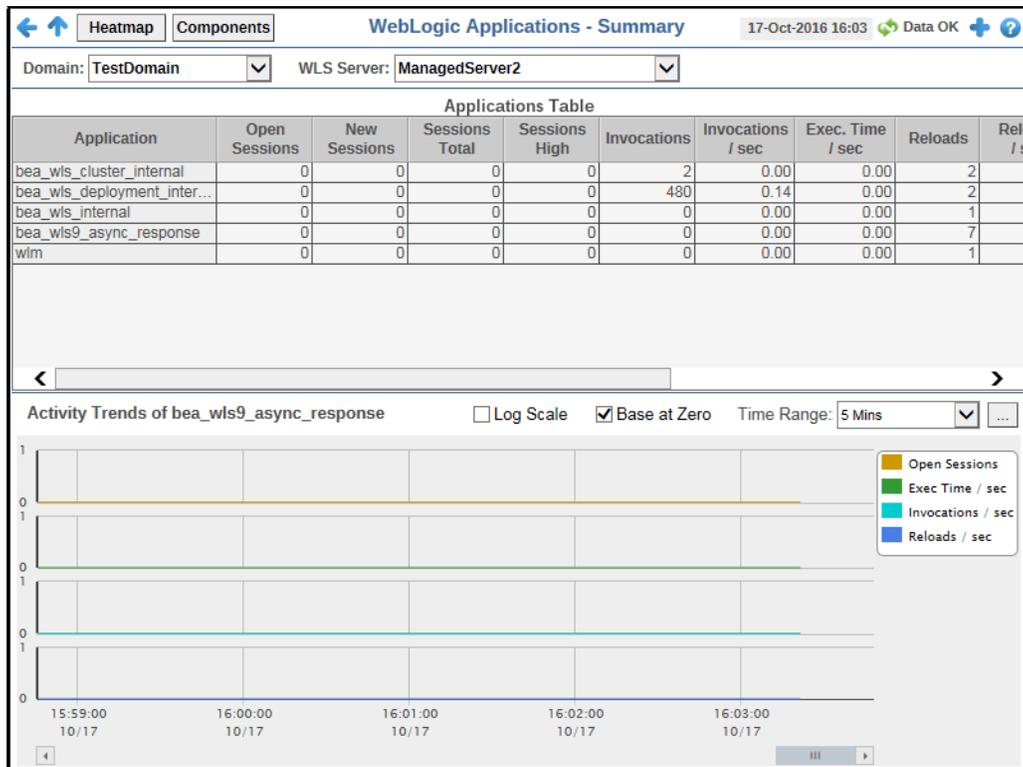
Select the metric driving the heatmap display. The default is **Alert Severity**. Each **Metric** has a color gradient bar that maps values to colors. The heatmap organizes the collections by connection, where each rectangle represents a collection. Mouse-over any rectangle to display the current values of the metrics for the collection. Click on a rectangle to drill-down to the associated ["Server Apps Summary"](#) display for a detailed view of metrics for that particular collection.

<b>Alert Severity</b>	<p>The maximum alert level in the item (index) associated with the rectangle. Values range from <b>0</b> to <b>2</b>, as indicated in the color gradient bar , where <b>2</b> is the greatest <b>Alert Severity</b>.</p> <p><b>2</b> -- Metrics that have exceeded their specified <b>ALARMLEVEL</b> threshold and have an Alert Severity value of <b>2</b> are shown in red. For a given rectangle, this indicates that one or more metrics have exceeded their alarm threshold.</p> <p><b>1</b> -- Metrics that have exceeded their specified <b>WARNINGLEVEL</b> threshold and have an Alert Severity value of <b>1</b> are shown in yellow. For a given rectangle, this indicates that one or more metrics have exceeded their warning threshold.</p> <p><b>0</b> -- Metrics that have not exceeded either specified threshold have an Alert Severity value of <b>0</b> and are shown in green. For a given rectangle, this indicates that no metrics have exceeded a specified alert threshold.</p>
<b>Alert Count</b>	<p>The total number of alarm and warning alerts in a given item (index) associated with the rectangle.</p> <p>The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>Open Sessions</b>	<p>The total number of open sessions in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from <b>0</b> to the alert threshold of <b>WlsOpenSessionsHigh</b>, which is <b>10</b>. The middle value in the gradient bar indicates the middle value of the range (the default is <b>5</b>).</p>

<b>Open Sessions/ sec</b>	The number of sessions opened per second in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of objects in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
<b>Invocations/sec</b>	The number of invocations per second in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of objects in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
<b>Invocation Total Count</b>	The total number of invocations in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of objects in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
<b>Exec Time/sec</b>	The execution time per second in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of objects in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
<b>Exec Time Total</b>	The total amount of execution time in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of objects in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
<b>Reload Total Count</b>	The total reload count in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of objects in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

## Server Apps Summary

Track performance, utilization, and trend data for all applications on a single WebLogic server.



**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu**, **Table** open commonly accessed displays.
- 6,047** The number of items currently in the display.
- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

**Note:** Fields with an asterisk (\*) at the end of the field definition contain data that is provided by the WebLogic Server MBean interface. Refer to WebLogic documentation for more information regarding these fields.

### Fields and Data

This display includes:

- Domain** Select the domain containing the WebLogic Server for which you want to view data.

**WLS Server**      Select the WebLogic server for which you want to view data.

### Applications Table

This table describes all topics on the selected server. Click a row to view metrics for a single topic in the “[Server Apps Trends](#)” display.

<b>Application</b>	The name of the application.
<b>Open Sessions</b>	The number of open sessions.
<b>New Sessions</b>	The number of new sessions.
<b>Sessions Total</b>	The total number of sessions.
<b>Invocations</b>	The number of invocations.
<b>Invocations/sec</b>	The number of invocations per second.
<b>Exec. Time/sec</b>	The rate of execution time in milliseconds per second.
<b>Reloads</b>	The number of reloads.
<b>Reloads / sec</b>	The rate of reloads (per second).
<b>Status</b>	The status of the application.
<b>Deployment State</b>	The current status of the application’s deployment.
<b>Expired</b>	This check box becomes automatically checked when the data displayed in the row has exceeded the specified cache expiration time (set by default at 45 seconds) and is no longer current. Once the cache has been refreshed and is displaying current data, the check box will return to being unchecked. This check box will remain unchecked as long as the cache has been refreshed within the specified cache expiration time and the data is current.
<b>time_stamp</b>	The date and time this row of data was last updated.

### Activity Trends of <application>

Shows the following:

**Open Sessions** -- Traces the total number of open sessions in the application.

**Exec Time/sec** -- Traces the execution time per second in the application.

**Invocations/sec** -- Traces the number of invocations per second.

**Reloads/sec** -- Traces the number of reloads per second.

### Log Scale

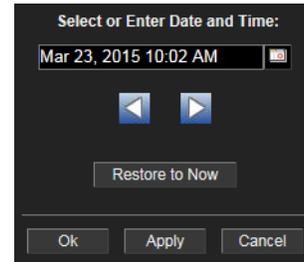
This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.

**Base at Zero**

When this option is checked, zero is set as the Y axis minimum for all graph traces.

**Time Range**

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



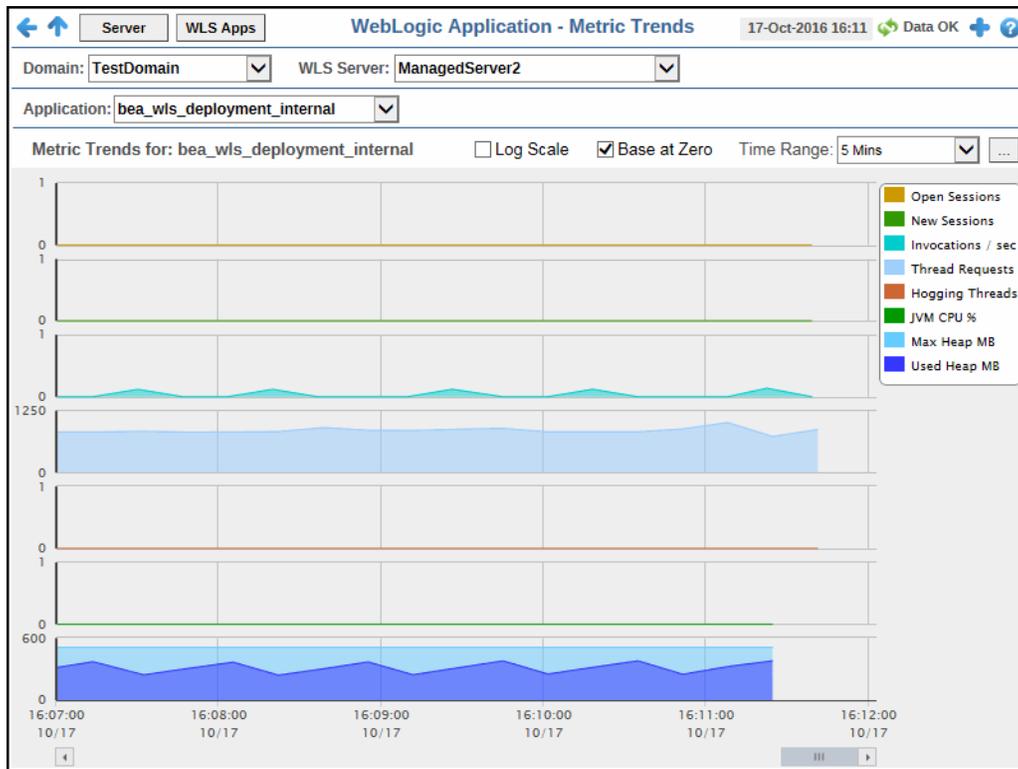
By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Server Apps Trends

View trend data for a single application on a particular WebLogic server.



**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

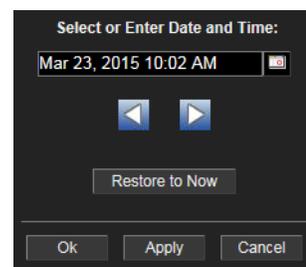
**Note:** Fields with an asterisk (\*) at the end of the field definition contain data that is provided by the WebLogic Server MBean interface. Refer to WebLogic documentation for more information regarding these fields.

### Fields and Data

This display includes:

- Domain** Select the domain containing the WebLogic Server for which you want to view data.

- WLS Server** Select the WebLogic server containing the application for which you want to view data.
- Application** Select the application for which you want to view data.
- Metric Trends for:** Shows message data for the selected collection.
- Open Sessions** -- Traces the total number of open sessions in the application.
  - New Sessions** -- Traces the number of new sessions in the application.
  - Invocations/sec** -- Traces the number of invocations per second in the application.
  - Thread Requests** -- Traces the number of thread requests in the application.
  - Hogging Threads** -- Traces the number of hogging threads in the application.
  - JVM CPU %** -- Traces the JVM CPU percentage in the application.
  - Max Heap MB** -- Traces the max heap used, in megabytes, in the application.
  - Used Heap MB** -- Traces the used heap, in megabytes, in the application.
- Log Scale** This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.
- Base at Zero** When this option is checked, zero is set as the Y axis minimum for all graph traces.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



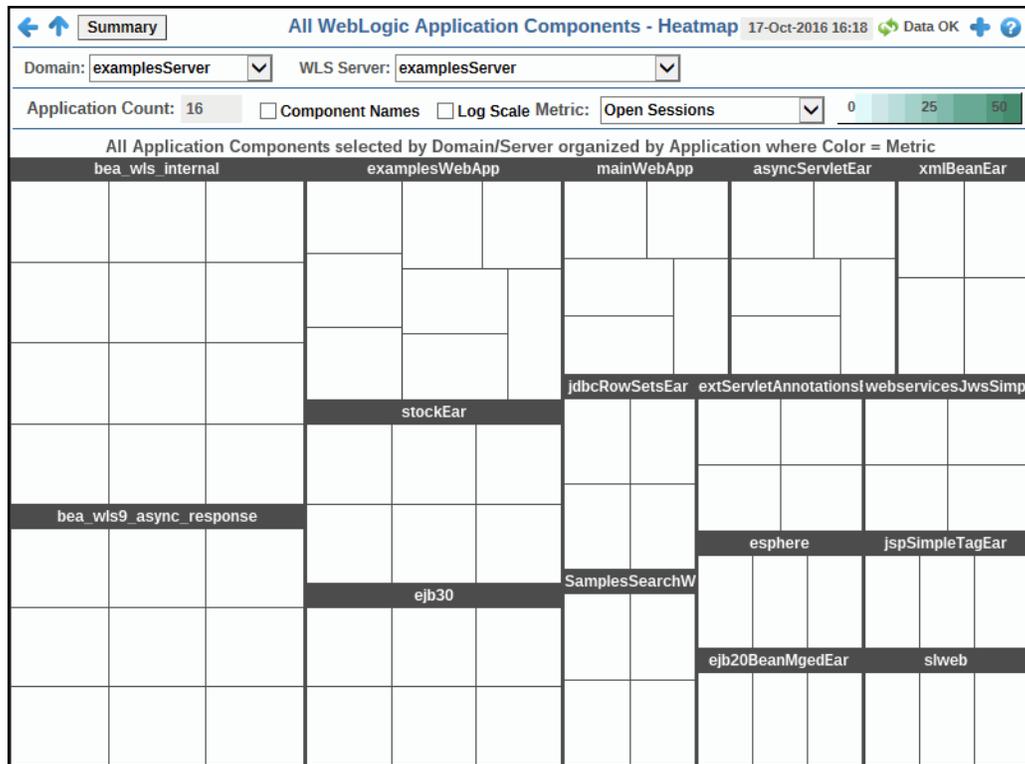
By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## App Components Heatmap

This display provides a heatmap view of the status and alerts of all application components contained within each application on a particular WebLogic server. The heatmap is organized so that each rectangle represents a collection contained within a specific connection. The rectangle color indicates the value of the selected metric in the **Metric** drop down list. You can mouse-over rectangles to view more details about the performance and status of each collection or click on a rectangle to drill-down to the “[WebLogic Single Application Summary](#)” display and view metrics for that particular collection. You can click the **Summary** button in this display to navigate to the “[Server Apps Summary](#)” display.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- The number of items currently in the display.

- Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

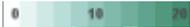
**Note:** Fields with an asterisk (\*) at the end of the field definition contain data that is provided by the WebLogic Server MBean interface. Refer to WebLogic documentation for more information regarding these fields.

## Fields and Data

This display includes:

- Domain** Select the domain containing the WebLogic Server for which you want to view data.
- WLS Server** Select the WebLogic server for which you want to view data.
- Application Count** The total number of application components in the WebLogic server.
- Component Names** Select this check box to display the names of the application components in the heatmap.
- Log Scale** This option enables visualization on a logarithmic scale, and should be used when the range in your data is very broad. For example, if you have data that ranges from the tens to the thousands, then data in the range of tens will be neglected visually if you do not check this option. This option makes data on both extreme ranges visible by using the logarithmic of the values rather than the actual values.

**Metric** Select the metric driving the heatmap display. The default is **Alert Severity**. Each **Metric** has a color gradient bar that maps values to colors. The heatmap organizes the collections by connection, where each rectangle represents a collection. Mouse-over any rectangle to display the current values of the metrics for the collection. Click on a rectangle to drill-down to the associated ["Server Apps Trends"](#) display for a detailed view of metrics for that particular collection.

**Open Sessions** The number of open sessions in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of objects in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

**Open Sessions/sec** The number of open sessions per second in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of objects in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

**Invocations/sec** The number of invocations per second in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of objects in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

**Invocation Total Count** The total number of invocations in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of objects in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

**Execution Time Total**

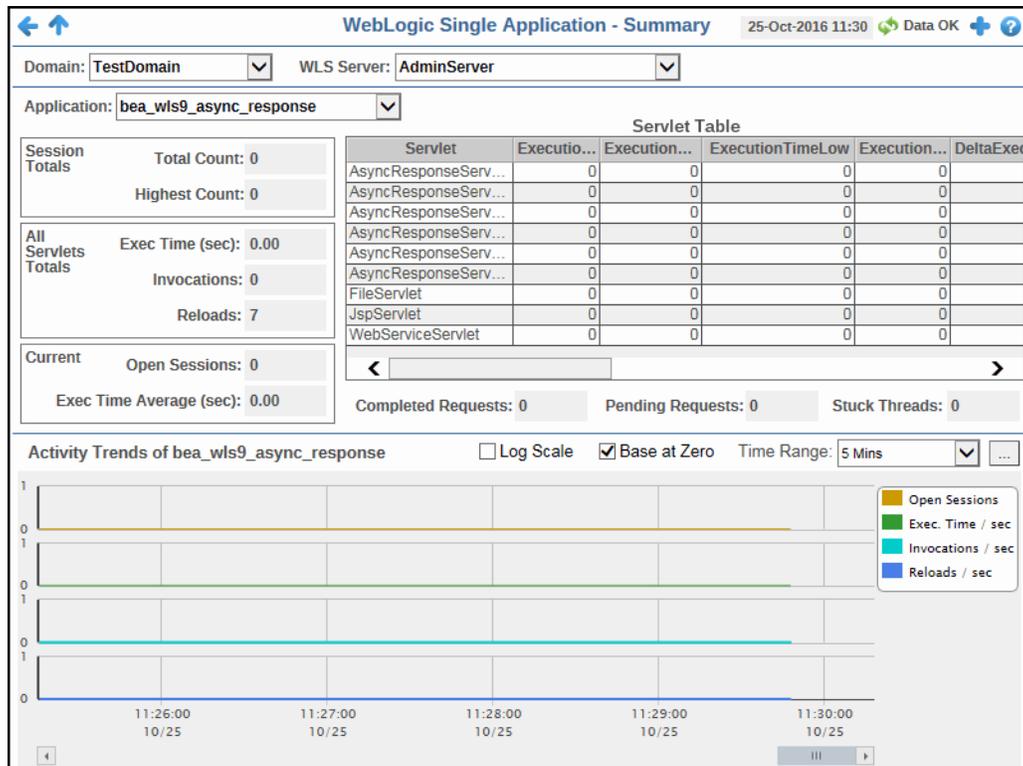
The total execution time in a given item (index) associated with the rectangle. The color gradient bar shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of objects in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

**Reload Total Count**

The total number of reloads in a given item (index) associated with the rectangle. The color gradient bar shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of objects in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

### WebLogic Single Application Summary

View performance, utilization, and trend data for a single application component on a single WebLogic server. This display is only accessible by clicking on an application component in the "App Components Heatmap".



**Title Bar** (possible features are):

-   Open the previous and upper display.
-  Open an instance of this display in a new window.
-  Open the online help page for this display.
- open commonly accessed displays.
- The number of items currently in the display.

 **Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

 Open the **Alert Views - RTView Alerts Table** display.

---

**Note:** Row columns in this table with an asterisk (\*) at the end of the column definition contain data that is provided by the WebLogic Server MBean interface. Refer to WebLogic documentation for more information regarding these columns.

---

### Fields and Data

This display includes:

- Domain** Select the domain containing the WebLogic server for which you want to view data.
- WLS Server** Select the WebLogic server containing the application for which you want to view data.
- Application** Select the application for which you want to view data.

#### Sessions Totals

- Total Count** The total number of sessions on the application.
- Highest Count** The highest total number of sessions on the application.

#### All Servlet Totals

- Exec Time (sec)** The amount of time all invocations of all servlets have executed since being created.\*
- Invocations** The total count of the times all servlets have been invoked.\*
- Reloads** The total count of the number of times all servlets have been reloaded.\*

#### Current

- Open Sessions** The number of currently open sessions.
- Exec Time Average (sec)** The current amount of time invocations of the servlet are being executed.\*

#### Servlet Table

This table describes metrics for each servlet on the application.

<b>Servlet</b>	The name of the servlet.
<b>Execution Time Average</b>	The average amount of time all invocations of the servlet have executed since being created.*
<b>Execution Time High</b>	The amount of time the single longest invocation of the servlet has executed since being created.*
<b>Execution Time Low</b>	The amount of time the single shortest invocation of the servlet has executed since being created.*
<b>Execution Time Total</b>	The total amount of time all invocations of the servlet have executed since being created.*
<b>Delta Execution Time</b>	The increase in the execution time (from the previous polling period to the current polling period).
<b>Rate Execution Time Total</b>	The total time taken to execute requests per second.
<b>Invocation Total Count</b>	The total count of the times this servlet has been invoked.*
<b>Delta Invocation Total Count</b>	The increase in the amount of invocations (from the previous polling period to the current polling period).
<b>Rate Invocation Total Count</b>	The total number of invocations per second.*
<b>Reload Total Count</b>	The total count of the number of times this servlet has been reloaded.*
<b>Delta Reload Total Count</b>	The increase in the amount of reloads (from the previous polling period to the current polling period).
<b>Rate Reload Total Count</b>	The number of times this servlet has been reloaded per second.
<b>Expired</b>	This check box becomes automatically checked when the data displayed in the row has exceeded the specified cache expiration time (set by default at 45 seconds) and is no longer current. Once the cache has been refreshed and is displaying current data, the check box will return to being unchecked. This check box will remain unchecked as long as the cache has been refreshed within the specified cache expiration time and the data is current.
<b>time_stamp</b>	The date and time this row of data was last updated.
<b>Completed Requests</b>	The total number of completed requests on the application.
<b>Pending Requests</b>	The total number of pending requests on the application.
<b>Stuck Threads</b>	The total number of stuck threads on the application.
<b>Metric Trends for: &lt;application&gt; Graph</b>	Shows message data for the selected collection. <ul style="list-style-type: none"> <li><b>Open Sessions</b> -- Traces the total number of open sessions in the application.</li> <li><b>Exec. Time/sec</b> -- Traces the number of executions per second in the application.</li> <li><b>Invocations/sec</b> -- Traces the number of invocations per second in the application.</li> <li><b>Reloads/sec</b> -- Traces the number of reloads per second in the application.</li> </ul>

**Log Scale**

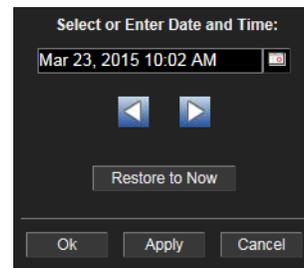
This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.

**Base at Zero**

When this option is checked, zero is set as the Y axis minimum for all graph traces.

**Time Range**

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



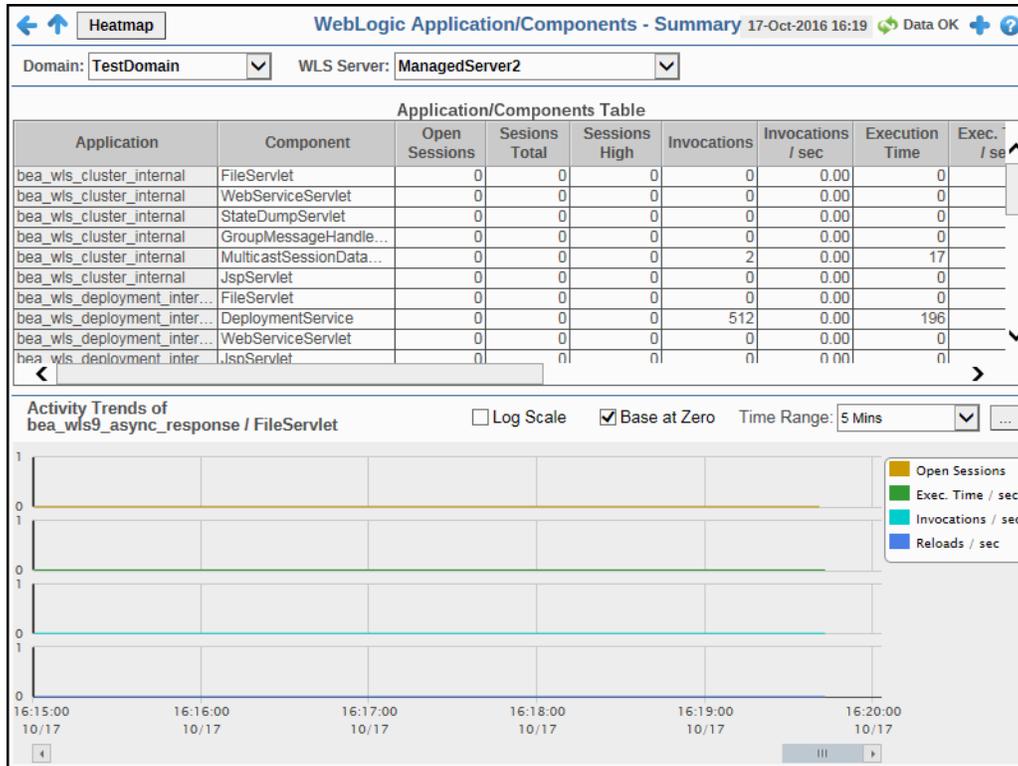
By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## App Components Summary

View performance, utilization, and trend data for all application components on a single WebLogic Server.



**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** , **Table** open commonly accessed displays.
- 6,047** The number of items currently in the display.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

Open the Alert Views - RTView Alerts Table display.

**Note:** Row columns in this table with an asterisk (\*) at the end of the column definition contain data that is provided by the WebLogic Server MBean interface. Refer to WebLogic documentation for more information regarding these columns.

### Fields and Data

This display includes:

- Domain** Select the domain containing the WebLogic server for which you want to view data.

**WLS Server** Select the WebLogic server containing the application for which you want to view data.

### Applications/Components Table

This table describes metrics for each application/component combination.

<b>Application</b>	The name of the application.
<b>Component</b>	The name of the component.
<b>Open Sessions</b>	The current total of open sessions.*
<b>Sessions Total</b>	The total number of sessions that have been opened.*
<b>Sessions High</b>	The highest number of sessions opened at one time.*
<b>Invocations</b>	The total number of invocations of the application or component.*
<b>Invocations/sec</b>	The rate of invocations (per second) of the application or component.*
<b>Execution Time</b>	The amount of time (in milliseconds) it took to execute the last invocation.*
<b>Exec. Time Low</b>	The average amount of time the single shortest invocation of the component has executed since it was most recently deployed.*
<b>Exec. Time High</b>	The average amount of time the single longest invocation of the component has executed since it was most recently deployed.**
<b>Exec. Time Average</b>	The average amount of time it took to execute all invocations of the component since it was most recently deployed.**
<b>Reloads</b>	The total number of times the WebLogic server has reloaded the component since it was last deployed.*
<b>Reloads/sec</b>	The rate of reloads (per second).*
<b>Status</b>	The status of the component.*
<b>Expired</b>	This check box becomes automatically checked when the data displayed in the row has exceeded the specified cache expiration time (set by default at 45 seconds) and is no longer current. Once the cache has been refreshed and is displaying current data, the check box will return to being unchecked. This check box will remain unchecked as long as the cache has been refreshed within the specified cache expiration time and the data is current.
<b>time_stamp</b>	The date and time this row of data was last updated.

### Metric Trends for: <application> Graph

Shows message data for the selected collection.

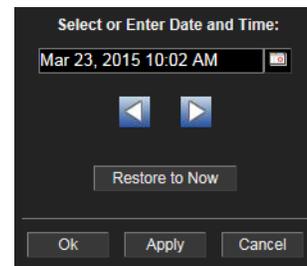
**Open Sessions** -- Traces the total number of open sessions in the application.

**Exec. Time/sec** -- Traces the number of executions per second in the application.

**Invocations/sec** -- Traces the number of invocations per second in the application.

**Reloads/sec** -- Traces the number of reloads per second in the application.

- Log Scale** This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.
- Base at Zero** When this option is checked, zero is set as the Y axis minimum for all graph traces.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows  to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

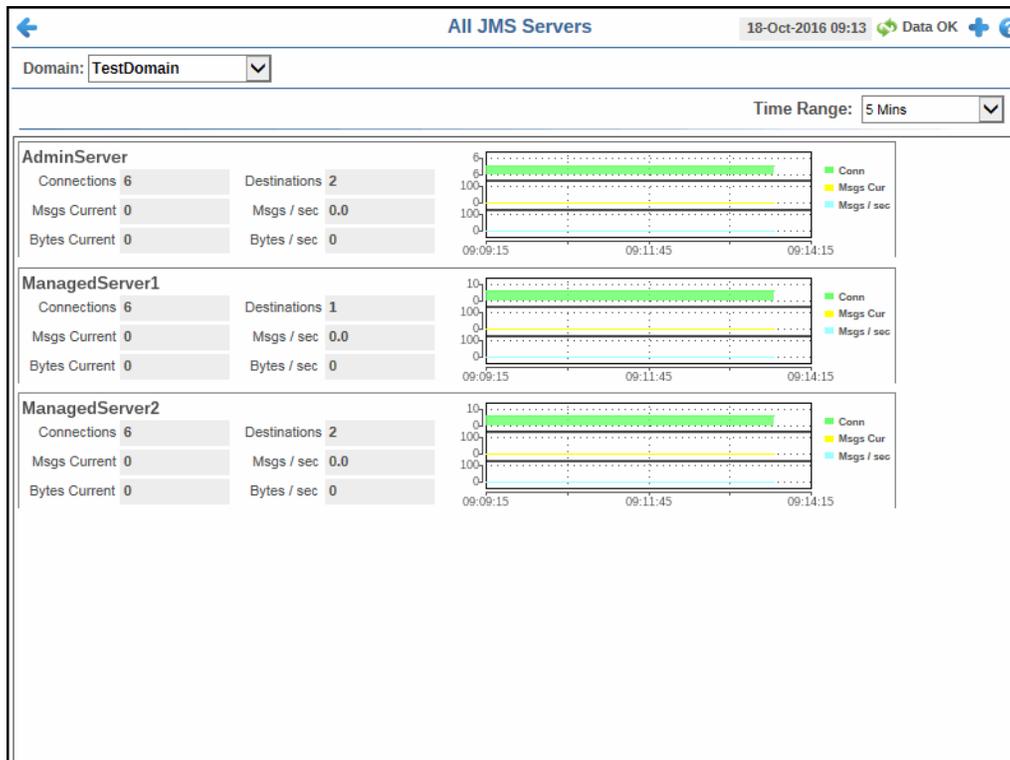
## JMS Servers View

These displays present several views of performance metrics for JMS Servers. This view contains the following:

- **"JMS Servers Grid"**: Track performance, utilization, and trend data for all JMS servers on a single domain.
- **"JMS Server Summary"**: View performance, utilization, and trend data for a particular JMS server.
- **"JMS Metric Trends"**: View activity trends for a particular JMS server.
- **"JMS Server Detail"**: Track performance and utilization metrics for all JMS servers on a single WebLogic server.

## JMS Servers Grid

Track performance, utilization, and trend data for all JMS servers on a single domain.



**Title Bar** (possible features are):

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu ▾, Table open commonly accessed displays.
- 6,047 The number of items currently in the display.
- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

**Note:** Row columns in this table with an asterisk (\*) at the end of the column definition contain data that is provided by the WebLogic Server MBean interface. Refer to WebLogic documentation for more information regarding these columns.

### Fields and Data

This display includes:

- Domain** Select the domain containing the WebLogic servers for which you want to view data.
- Time Range** Select the period of time from this drop down for which you want to view data.

**Server Trend Graphs**

Trend graphs and the following metrics displays for servers within the selected domain.

**Conn** -- Traces the total number of connections on the server.

**Msgs Cur** -- Traces the current number of incoming/outgoing messages on the server.

**Msgs/sec** -- Traces the number of incoming/outgoing messages per second on the server.

**Connections** The number of the connections on the server.

**Msgs Current** The current number of messages stored on the JMS server. This number does not include the pending messages.\*

**Bytes Current** The current number of bytes stored on the JMS server.\*

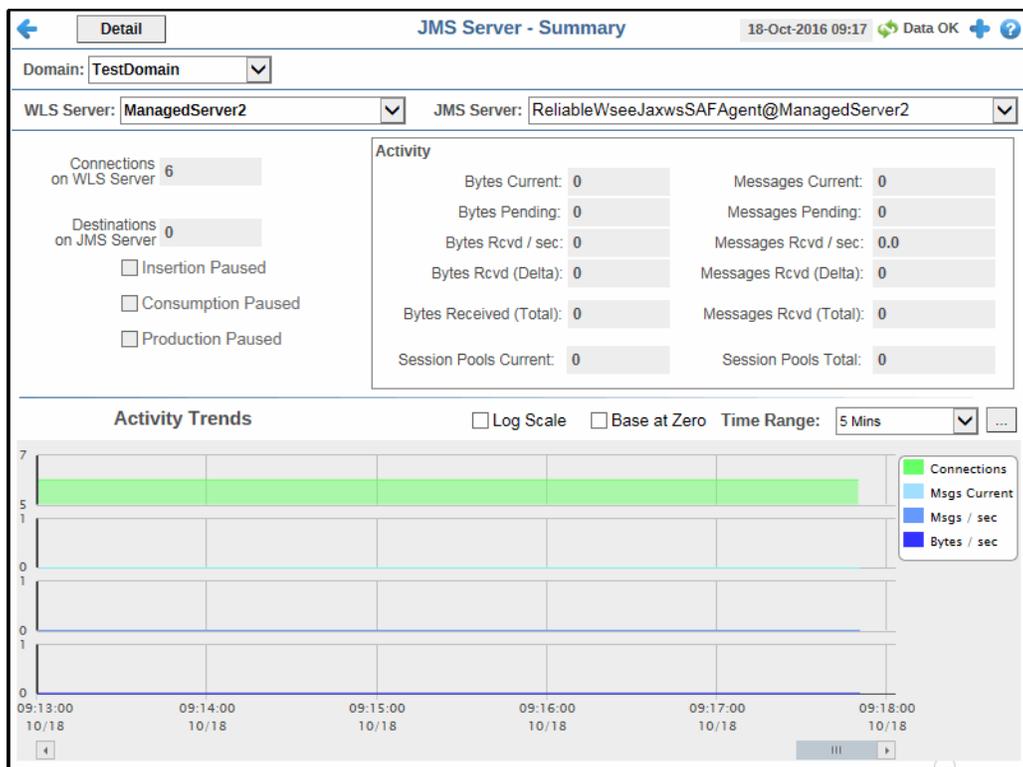
**Destinations** The current number of destinations stored on the JMS server.\*

**Msgs/sec** The rate of messages received (per second) on the JMS server.\*

**Bytes/sec** The rate of bytes received (per second) on the JMS server.\*

**JMS Server Summary**

View performance, utilization, and trend data for a particular JMS server.



**Title Bar** (possible features are):

-   Open the previous and upper display.
-  Open an instance of this display in a new window.
-  Open the online help page for this display.
- open commonly accessed displays.
- The number of items currently in the display.

 **Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

 Open the **Alert Views - RTView Alerts Table** display.

**Note:** Row columns in this table with an asterisk (\*) at the end of the column definition contain data that is provided by the WebLogic Server MBean interface. Refer to WebLogic documentation for more information regarding these columns.

### Fields and Data

This display includes:

<b>Domain</b>	Select the domain containing the WebLogic server for which you want to view data.
<b>WLS Server</b>	Select the WebLogic server containing the JMS server for which you want to view data.
<b>JMS Server</b>	Select the JMS server for which you want to view data.
<b>Connections on WLS Server</b>	The total number of connections on the WebLogic server.
<b>Destinations on JMS Server</b>	The number of destinations on the selected JMS server.
<b>Insertion Paused</b>	The current insertion paused state of the JMS server.*
<b>Consumption Paused</b>	The current consumption paused state of the JMS server.*
<b>Production Paused</b>	The current production paused state of the JMS server.*
<b>Activity Region</b>	
<b>Bytes Current</b>	The current number of bytes stored on the JMS server.*
<b>Bytes Pending</b>	The current number of bytes pending (unacknowledged or uncommitted) on the JMS server.*
<b>Bytes Rcvd/sec</b>	The number of bytes received per second.
<b>Bytes Rcvd (Delta)</b>	The increase in the amount of bytes received (from the previous polling period to the current polling period).*
<b>Bytes Rcvd (Total)</b>	The total number of bytes received since the JMS server was last restarted.*

<b>Alert Session Pools Current</b>	The current number of session pools instantiated on the server.*
<b>Messages Current</b>	The current number of messages stored on the JMS server.*
<b>Messages Pending</b>	The current number of messages pending (unacknowledged or uncommitted) on the JMS server.*
<b>Messages Rcvd/sec</b>	The number of messages received per second since the server was last restarted.*
<b>Messages Rcvd (Delta)</b>	The increase in the amount of messages received (from the previous polling period to the current polling period).
<b>Messages Rcvd (Total)</b>	The total number of messages received since the server was last restarted.*
<b>Session Pools Total</b>	The total number of session pools on the server.*

### Activity Trends

Shows message data for the selected collection.

**Connections** -- Traces the total number of connections on the server.

**Msgs Current** -- Traces the number of current messages.

**Msgs/sec** -- Traces the number of incoming/outgoing messages per second.

**Bytes/sec** -- Traces the number incoming/outgoing bytes per second.

### Log Scale

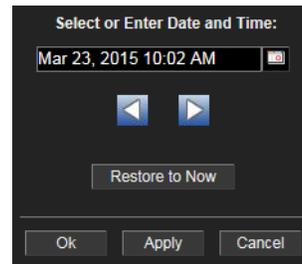
This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.

**Base at Zero**

When this option is checked, zero is set as the Y axis minimum for all graph traces.

**Time Range**

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



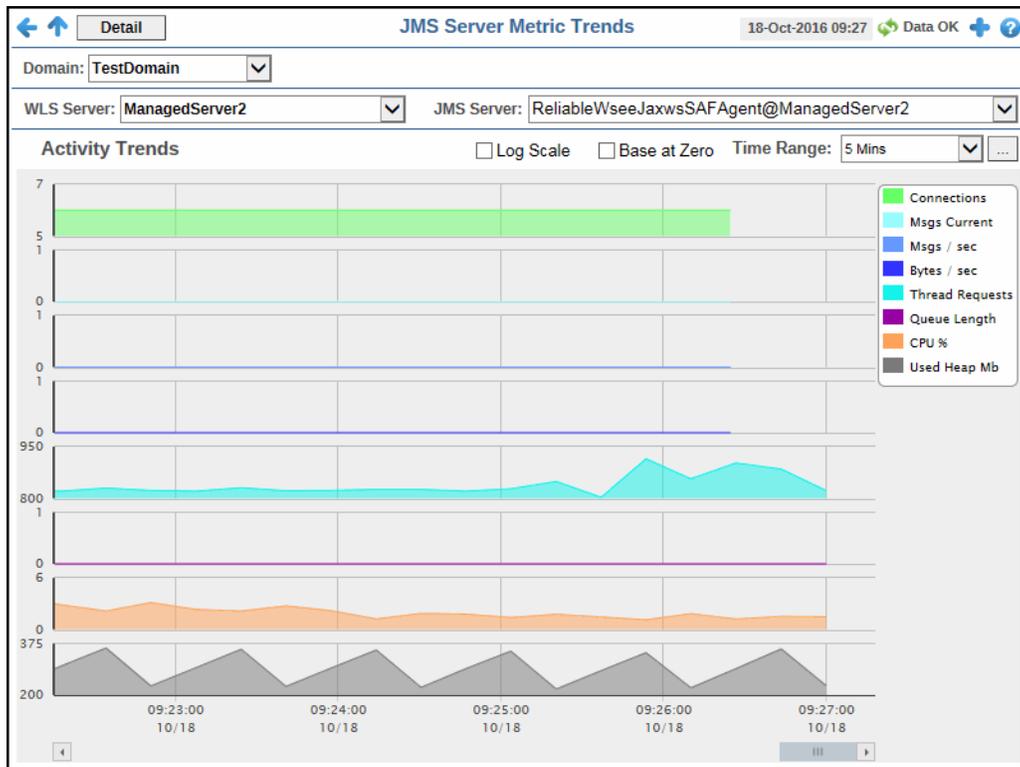
By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## JMS Metric Trends

View activity trends for a particular JMS server.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

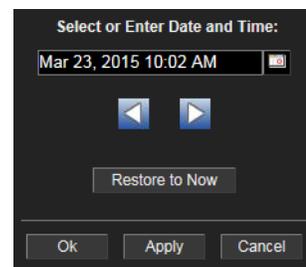
**Note:** Fields with an asterisk (\*) at the end of the field definition contain data that is provided by the WebLogic Server MBean interface. Refer to WebLogic documentation for more information regarding these fields.

### Fields and Data

This display includes:

- Domain** Select the domain containing the WebLogic server for which you want to view data.

- WLS Server** Select the WebLogic server containing the JMS server for which you want to view data.
- JMS Server** Select the JMS server for which you want to view data.
- Activity Trends** Shows message data for the selected collection.
- Connections** -- Traces the total number of connections on the server.
  - Msgs Current** -- Traces the number of current messages stored on the server.
  - Msgs/sec** -- Traces the number of messages stored per second.
  - Bytes/sec** -- Traces the number bytes received on the server per second.
  - Thread Requests** -- Traces the number of thread requests.
  - Queue Length** -- Traces the length of the queue.
  - CPU %** -- Traces the percentage of CPU used.
  - Used Heap Mb** -- Traces the amount of heap used, in megabytes.
- Log Scale** This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.
- Base at Zero** When this option is checked, zero is set as the Y axis minimum for all graph traces.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## JMS Server Detail

View performance and utilization metrics for all JMS servers on a single WebLogic server.

The screenshot shows the 'JMS Servers - Detail Tables' interface. At the top, there is a 'Summary' tab and a title bar with the text 'JMS Servers - Detail Tables'. The interface includes several filters: 'Domain: TestDomain', 'DataServer: WLM-alpha-SLHOST-WIN7', 'WLS Server: ManagedServer2', and 'Cluster: TestCluster-114'. The main content area is divided into two tables.

**JMS Server Runtime Table:**

Name	Destinations Current	Destinations Total	Messages Received	Msgs Rcvd Delta	Msgs Rcvd per sec	Messages Current	Mess Pen
ReliableWseeJaxwsSAFAgent@Ma...	0	0	0	0	0.00	0	
ReliableWseeSAFAgent@Managed...	0	0	0	0	0.00	0	
WseeJaxwsJmsServer_auto_2	4	4	0	0	0.00	0	
WseeJmsServer_auto_3	2	2	0	0	0.00	0	

**JMS Runtime Table:**

Connections Current	Connections High	Connections Total	Connections Delta	JMS Servers Current	JMS Servers High	JMS Servers Total	
6	6	6	0	4	4	4	Compd

**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

**Note:** Row columns in this table with an asterisk (\*) at the end of the column definition contain data that is provided by the WebLogic Server MBean interface. Refer to WebLogic documentation for more information regarding these columns.

### Fields and Data

This display includes:

- Domain** Select the domain containing the WebLogic server for which you want to view data.
- DataServer** The name of the data server associated with the domain.

**WLS Server** Select the WebLogic server containing the JMS server for which you want to view data.

**Cluster** The name of the cluster associated with the WebLogic Server.

### JMS Server Runtime Table

This table describes all runtime statistics on the selected server. Click a row to view metrics for a single server in the ["JMS Server Summary"](#) display.

<b>Name</b>	The name of the JMS Server.
<b>Destinations Current</b>	The current number of destinations for the JMS server.*
<b>Destinations Total</b>	The total number of destinations instantiated on the JMS server since the last polling update.*
<b>Messages Received</b>	The total number of messages received by the JMS server since the last polling update.*
<b>Msgs Rcvd Delta</b>	The increase in the amount of messages received (from the previous polling period to the current polling period).
<b>Msgs Rcvd per sec</b>	The rate of messages received (per second) by the server.
<b>Messages Current</b>	The current number of messages stored on the JMS server.*
<b>Messages Pending</b>	The current number of pending messages stored on the JMS server.*
<b>Messages High</b>	The highest number of total messages stored by the JMS server since the last polling update.*
<b>Bytes Received</b>	The number of bytes received by the server since the last polling update.
<b>Bytes Rcvd Delta</b>	The increase in the amount of bytes received (from the previous polling period to the current polling period).
<b>Bytes Rcvd per sec</b>	The rate of bytes received (per second) by the server.*
<b>Bytes Current</b>	The current number of bytes stored on the JMS server.*
<b>Bytes Pending</b>	The current number of bytes pending that are stored on the JMS server.*
<b>Bytes High</b>	The largest number of bytes stored on the JMS server since the last polling update.*
<b>Health State</b>	The current health state of the JMS server.*
<b>Consumption Paused</b>	The current consumption paused state of the JMS server. When checked, consumption is paused.*
<b>Cons Paused Enabled</b>	The current consumption paused state on the JMS server (enabled/disabled).*
<b>Insertion Paused</b>	The current insertion paused state of the JMS server.*
<b>Ins Paused State</b>	The current insertion paused state of the JMS server (enabled/disabled).*

<b>Production Paused</b>	The current production paused state of the JMS server.*
<b>Prod Paused State</b>	The current production paused state of the JMS server (enabled/disabled).*
<b>Session Pools Current</b>	The current number of session pools instantiated on the JMS server.*
<b>Session Pools High</b>	The highest number of session pools instantiated on the server since the last polling update.*
<b>Session Pools Total</b>	The total number of session pools instantiated on the JMS server since the last polling update.*
<b>Destinations</b>	The number of destinations instantiated on the JMS server since the last polling interval.
<b>Transactions</b>	The number of transactions that exist on the JMS server.*
<b>Transactions Pending</b>	The number of pending transactions that exist on the JMS server.*
<b>Server Runtime</b>	The name of the JMS Server.*
<b>time_stamp</b>	The date and time this row of data was last updated.

### JMS Runtime Table

This table describes performance metrics about the connectivity of the JMS Server.

<b>Connections Current</b>	The current number of connections to the WebLogic server.*
<b>Connections High</b>	The highest number of connections made to the WebLogic server since the last polling update.*
<b>Connections Total</b>	The total number of connections made to the WebLogic server since the last polling update.*
<b>Connections Delta</b>	The increase in the amount of connections (from the previous polling period to the current polling period).
<b>JMS Servers Current</b>	The number of JMS servers currently deployed on the WebLogic server.*
<b>JMS Servers High</b>	The highest number of JMS servers that were deployed on the WebLogic server since the server was started.*
<b>JMS Servers Total</b>	The total number of JMS servers that were deployed on the WebLogic server since the server was started.*
<b>HealthState</b>	The current state of health of the JMS service.*
<b>JMSServers</b>	The number of JMS servers deployed on the WebLogic server.*
<b>Connections</b>	The number of connections to the WebLogic server.*
<b>JMSPooledConnections</b>	The number of JMS pooled connections on the server.
<b>ServerRuntime</b>	The name of the server.
<b>time_stamp</b>	The date and time this row of data was last updated.

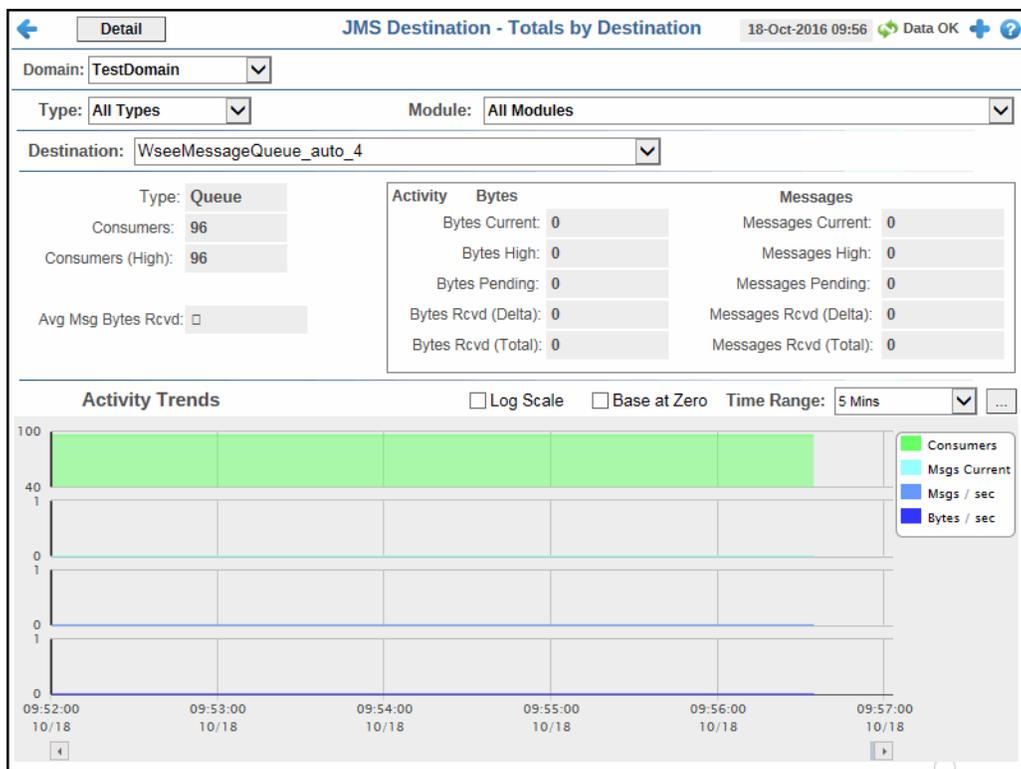
## JMS Destinations View

These displays present several views of performance metrics for JMS Destinations. This view contains the following:

- **"Totals by Destination"**: View performance, utilization, and trend data for a particular destination.
- **"Detail by Destination"**: A tabular view that allows you to view performance and utilization metrics for destinations for each JMS Server, as well as destination metrics across all JMS Servers.
- **"Distribution by Server"**: Shows performance and utilization metrics and trends for a destination on a single domain.
- **"Destination by Server"**: Shows metrics and trends for a particular destination on a JMS Server.
- **"Detail by Server"**: Shows metrics for all destinations on one or all modules on a JMS Server.

### Totals by Destination

View performance, utilization, and trend data for a particular destination.



**Title Bar** (possible features are):

-   Open the previous and upper display.
-  Open an instance of this display in a new window.
-  Open the online help page for this display.
- ,  open commonly accessed displays.
- The number of items currently in the display.

 **Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

 Open the **Alert Views - RTView Alerts Table** display.

---

**Note:** Row columns in this table with an asterisk (\*) at the end of the column definition contain data that is provided by the WebLogic Server MBean interface. Refer to WebLogic documentation for more information regarding these columns.

---

### Fields and Data

This display includes:

<b>Domain</b>	Select the domain containing the WebLogic server for which you want to view data.
<b>Type</b>	Select the type of destination ( <b>Queue</b> or <b>Topic</b> ) for which you want to view data, or select <b>All Types</b> .
<b>Module</b>	Select the module containing the destination for which you want to view data.
<b>Destination</b>	Select the destination for which you want to view data.
<b>Type</b>	The type of destination (queue or topic) selected.
<b>Consumers</b>	The current number of consumers accessing the destination.
<b>Consumers (High)</b>	The highest number of consumers accessing the destination since the last polling update.
<b>Avg Msg Bytes Rcvd</b>	The average number of bytes received in the destination since the last polling update.

### Activity Region

#### Bytes

<b>Bytes Current</b>	The current number of bytes stored in the destination.*
<b>Bytes High</b>	The highest number of bytes stored in the destination since the last polling update.*
<b>Bytes Pending</b>	The current number of pending bytes stored in the destination.*
<b>Bytes Rcvd (Delta)</b>	The increase in the amount of bytes received (from the previous polling period to the current polling period).
<b>Bytes Rcvd (Total)</b>	The number of bytes received in this destination since the last polling update.*

#### Messages

<b>Messages Current</b>	The current number of messages in the destination.*
<b>Messages High</b>	The highest number of messages in the destination since the last polling update.*
<b>Messages Pending</b>	The current number of pending messages in the destination.*
<b>Messages Rcvd (Delta)</b>	The increase in the amount of messages received (from the previous polling period to the current polling period).
<b>Messages Rcvd (Total)</b>	The total number of messages in the destination since the last polling update.*

**Activity Trends**

Shows message data for the selected collection.

**Consumers** -- Traces the total number of consumers on the destination.

**Msgs Current** -- Traces the number of current messages.

**Msgs/sec** -- Traces the number of messages received per second in the destination.

**Bytes/sec** -- Traces the number of bytes received per second in the destination.

**Log Scale**

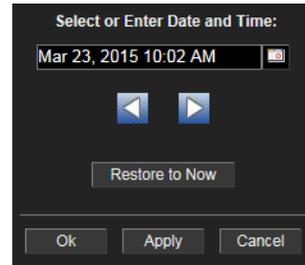
This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.

**Base at Zero**

When this option is checked, zero is set as the Y axis minimum for all graph traces.

**Time Range**

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Detail by Destination

View destination statistics by individual JMS Servers and across all servers.

**JMS Destinations - Detail by Destination** 18-Oct-2016 10:00 Data OK

Domain: TestDomain  
 Type: All Types Module: All Modules  
 Destination: WseeMessageQueue\_auto\_4

Destination Statistics for each JMS Server

JMS Server	Type	Bytes Current	Bytes High	Bytes Pending	Bytes Rcvd	Rcvd Delta	Rcvd Rate	Consumers Current
WseeJmsServer_auto_3	Queue	0	0	0	0	0	0.0	

Totals across All Servers

Bytes Current	Bytes High	Bytes Pending	Bytes Rcvd	Rcvd Delta	Rcvd Rate	Consumers Current	Consumers High	Consumers Total	Msg Current
0	0	0	0	0	0.0	96	96	96	

**Title Bar** (possible features are):

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu Table open commonly accessed displays.
- 6,047 The number of items currently in the display.
- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

**Note:** Row columns in this table with an asterisk (\*) at the end of the column definition contain data that is provided by the WebLogic Server MBean interface. Refer to WebLogic documentation for more information regarding these columns.

### Fields and Data

This display includes:

- Domain** Select the domain containing the destination for which you want to view data.
- Type** Select the type of destination (queue or topic) for which you want to view data, or select **All Types**.

**Module** Select the module containing the destination for which you want to view data.

**Destination** Select the destination for which you want to view data.

### Destination Statistics for each JMS Server Table

Each row in this table describes all destination statistics for a particular JMS Server. Click a row to view additional metrics in the ["Destination by Server"](#) display.

<b>JMS Server</b>	The name of the JMS Server.
<b>Type</b>	The type of destination (queue or topic).
<b>Bytes Current</b>	The current number of bytes stored in the destination.*
<b>Bytes High</b>	The highest number of bytes stored in the destination since the last polling update.*
<b>Bytes Pending</b>	The current number of pending bytes stored in the destination.*
<b>Bytes Rcvd</b>	The number of bytes received in this destination since the last polling update.*
<b>Rcvd Delta</b>	The increase in the amount of bytes received (from the previous polling period to the current polling period).
<b>Rcvd Rate</b>	The rate of bytes received (per second) into the destination.*
<b>Consumers Current</b>	The current number of consumers accessing the destination.*
<b>Consumers High</b>	The highest number of consumers accessing the destination since the last polling update.*
<b>Consumers Total</b>	The total number of consumers accessing the destination since the last polling update.*
<b>Msgs Current</b>	The current number of messages in the destination.*
<b>Msgs Deleted</b>	The number of messages that have been deleted from the destination.*
<b>Msgs High</b>	The highest number of messages in the destination since the last polling update.*
<b>Msgs Moved</b>	The number of moved messages in the destination since the last polling update.*
<b>Msgs Pending</b>	The current number of pending messages in the destination.*
<b>Msgs Rcvd</b>	The number of messages received in the destination since the last polling update.*
<b>Rcvd Delta</b>	The increase in the amount of messages received (from the previous polling period to the current polling period).
<b>Rcvd Rate</b>	The rate of messages received (per second) in the destination since the last polling update.*
<b>Msgs Thresh Time</b>	The amount of time in the threshold condition since the last polling update.

<b>Bytes Thresh Time</b>	The amount of time in the threshold condition since the last polling update.
<b>Paused</b>	Indicates whether or not the destination is paused at the current time.*
<b>Consumption Paused</b>	Indicates the consumption paused state of the destination.
<b>Consumption Paused State</b>	The current consumption paused state of the destination.
<b>Production Paused</b>	Indicates the production paused state of the destination.
<b>Production Paused State</b>	The current production paused state of the destination.
<b>Insertion Paused</b>	Indicates the insertion paused state of the destination.
<b>Insertion Paused state</b>	The current insertion pause state of the destination.
<b>State</b>	The current health state of the destination.
<b>time_stamp</b>	The date and time this row of data was last updated.

#### Totals across All Servers Table

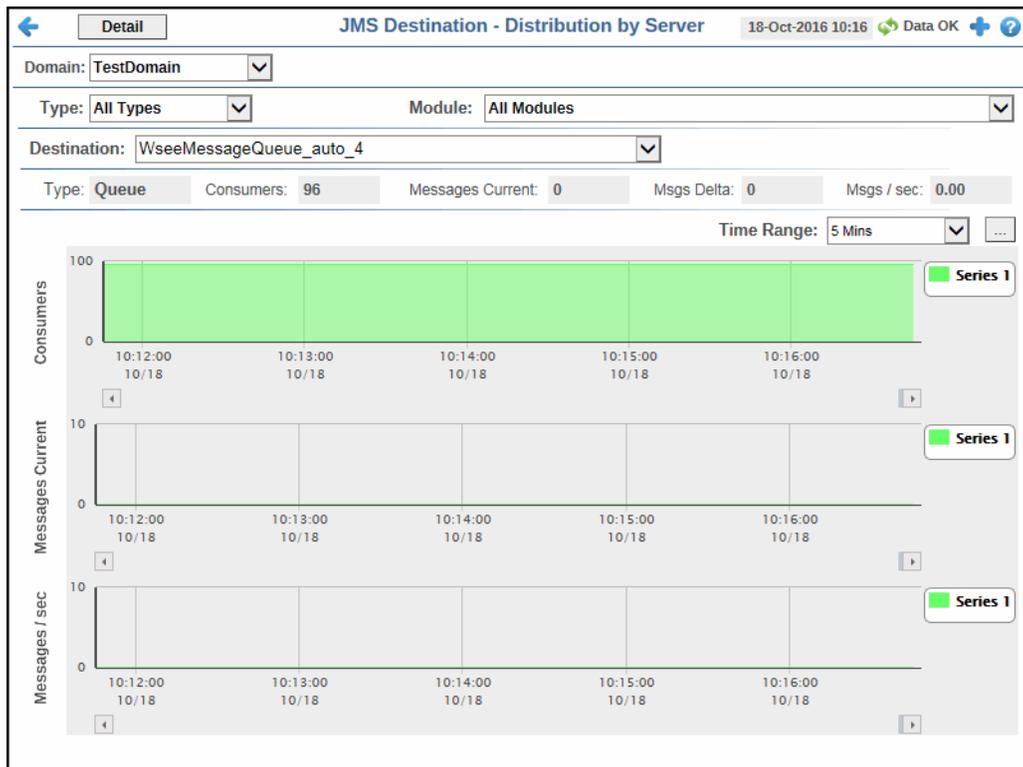
This table describes destination metric totals across all servers.

<b>Bytes Current</b>	The current number of bytes stored across all servers.
<b>Bytes High</b>	The highest number of bytes on any one of the servers at one time since the last polling update.*
<b>Bytes Pending</b>	The total number of bytes pending on all servers.*
<b>Bytes Rcvd</b>	The number of bytes received across all servers.*
<b>Rcvd Delta</b>	The increase in the amount of bytes received (from the previous polling period to the current polling period).
<b>Rcvd Rate</b>	The rate of messages received (per second) by all servers.
<b>Consumers Current</b>	The number of current consumers across all servers.
<b>Consumers High</b>	The highest number of consumers at one time across all servers.*
<b>Consumers Total</b>	The total number of consumers across all servers.*
<b>Msgs Current</b>	The current number of messages stored across all destinations on the all servers.*
<b>Msgs Deleted</b>	The number of deleted messages across all destinations on all servers.*
<b>Msgs High</b>	The highest number of total messages stored across all destinations for all servers.*
<b>Msgs Moved</b>	The number of messages across all destinations for all servers that were moved.*
<b>Msgs Pending</b>	The current number of pending messages across all servers.

<b>Msgs Rcvd</b>	The total number of messages received across all servers.
<b>Rcvd Delta</b>	The increase in the amount of messages received (from the previous polling period to the current polling period).
<b>Rcvd Rate</b>	The rate of messages received (per second) across all servers.

## Distribution by Server

Track performance and utilization metrics and trends for a destination on a single domain.



**Title Bar** (possible features are):

-   Open the previous and upper display.
-  Open an instance of this display in a new window.
-  Open the online help page for this display.
- ,  open commonly accessed displays.
- The number of items currently in the display.

 **Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

 Open the **Alert Views - RTView Alerts Table** display.

---

**Note:** Fields with an asterisk (\*) at the end of the field definition contain data that is provided by the WebLogic Server MBean interface. Refer to WebLogic documentation for more information regarding these fields.

---

### Fields and Data

This display includes:

<b>Domain</b>	Select the domain containing the destination for which you want to view data.
<b>Type</b>	Select the type of destination (queue or topic) for which you want to view data, or select <b>All Types</b> .
<b>Module</b>	Select the module containing the destination for which you want to view data.
<b>Destination</b>	Select the destination for which you want to view data.
<b>Type</b>	The type of destination (queue or topic).
<b>Consumers</b>	The number of consumers.
<b>Msgs Current</b>	The current number of messages on the destination.*
<b>Msgs Delta</b>	The increase in the amount of messages received (from the previous polling period to the current polling period).
<b>Msgs/sec</b>	The rate of messages received (per second) by the destination.

**Server Trend Graphs**

Trend graphs and the following metrics displays for servers within the selected domain.

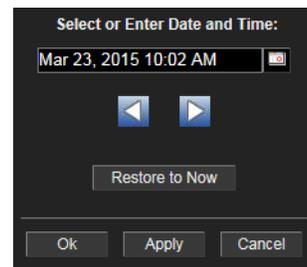
**Consumers** -- Traces the total number of consumers on the server.

**Messages Current** -- Traces the current number of incoming/outgoing messages on the server.

**Messages/sec** -- Traces the number of incoming/outgoing messages per second on the server.

**Time Range**

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



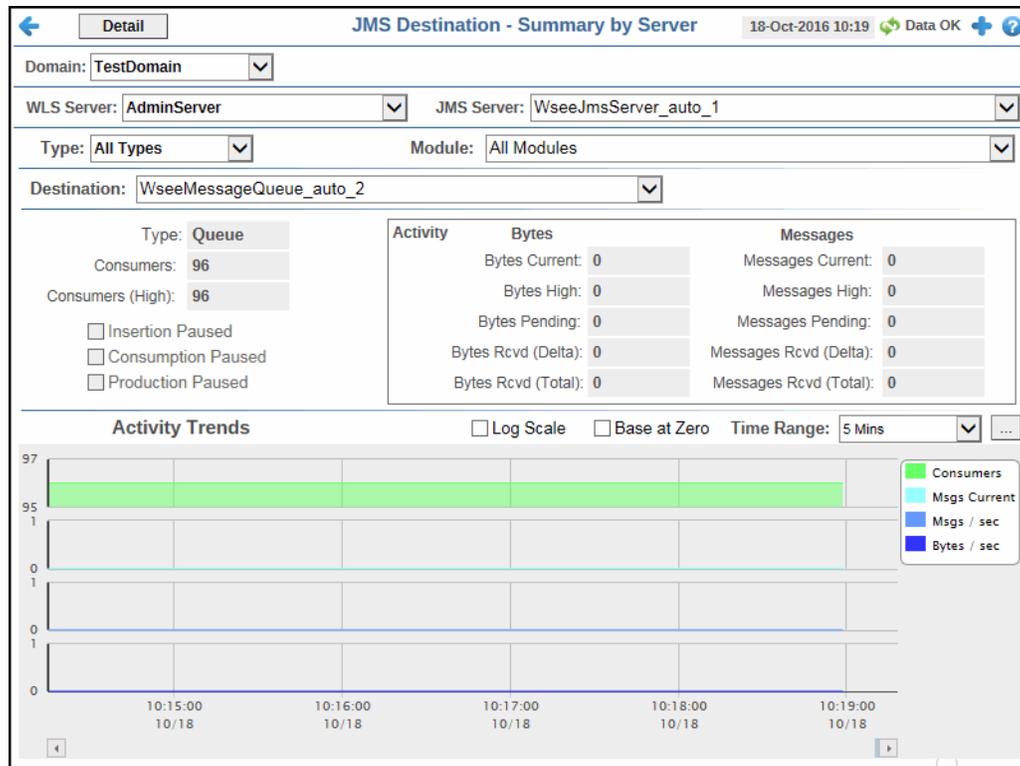
By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Destination by Server

Track performance, utilization, and trend data for a particular destination on a single JMS Server.



**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.
- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

**Note:** Fields with an asterisk (\*) at the end of the field definition contain data that is provided by the WebLogic Server MBean interface. Refer to WebLogic documentation for more information regarding these fields.

### Fields and Data

This display includes:

- Domain** Select the domain containing the WebLogic Server for which you want to view data.

<b>WLS Server</b>	Select the WebLogic server containing the JMS Server for which you want to view data.
<b>JMS Server</b>	Select the JMS Server containing the destination for which you want to view data
<b>Type</b>	Select the type of destination (queue or topic) for which you want to view data, or select <b>All Types</b> .
<b>Module</b>	Select the module containing the destination for which you want to view data.
<b>Destination</b>	Select the destination for which you want to view data.
<b>Type</b>	The type (queue or topic) of the selected destination.
<b>Consumers</b>	The current number of consumers accessing the destination.*
<b>Consumer (High)</b>	The highest number of consumers accessing the destination since the last polling update.*
<b>Insertion Paused</b>	Indicates the insertion paused state of the destination.*
<b>Consumption Paused</b>	Indicates the consumption paused state of the destination.*
<b>Production Paused</b>	Indicates the production paused state of the destination.*
<b>Activity Region</b>	
<b>Bytes Current</b>	The current number of bytes stored in the destination.*
<b>Bytes High</b>	The highest number of bytes stored in the destination since the last polling update.*
<b>Bytes Pending</b>	The current number of pending bytes stored in the destination.*
<b>Bytes Rcvd (Delta)</b>	The increase in the amount of bytes received (from the previous polling period to the current polling period).
<b>Bytes Rcvd (Total)</b>	The number of bytes received in this destination since the last polling update.*
<b>Messages Current</b>	The current number of messages in the destination.*
<b>Messages High</b>	The highest number of messages in the destination since the last polling update.*
<b>Messages Pending</b>	The current number of pending messages in the destination.*
<b>Messages Rcvd (Delta)</b>	The increase in the amount of messages received (from the previous polling period to the current polling period).
<b>Messages Rcvd (Total)</b>	The number of messages received in the destination since the last polling update.*
<b>Activity Trends</b>	Shows message data for the selected collection. <b>Consumers</b> -- Traces the total number of consumers accessing the destination. <b>Msgs Current</b> -- Traces the number of current messages. <b>Msgs/sec</b> -- Traces the number of messages received per second. <b>Bytes/sec</b> -- Traces the number bytes received per second.

**Log Scale**

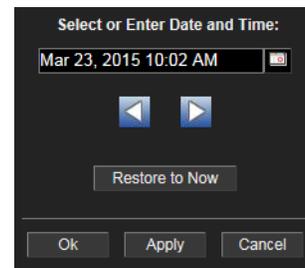
This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.

**Base at Zero**

When this option is checked, zero is set as the Y axis minimum for all graph traces.

**Time Range**

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Detail by Server

Track performance and utilization metrics for all destinations on a single JMS Server.

Destination Name	Type	Bytes Current	Bytes High	Bytes Pending	Bytes Rcvd	Rcvd Delta	Rcvd Rate
WseeBufferedRequestErrorQueue_auto_2	Queue	0	0	0	0	0	0.0
WseeBufferedRequestQueue_auto_2	Queue	0	0	0	0	0	0.0
WseeBufferedResponseErrorQueue_auto_2	Queue	0	0	0	0	0	0.0
WseeBufferedResponseQueue_auto_2	Queue	0	0	0	0	0	0.0

### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

**Note:** Row columns in this table with an asterisk (\*) at the end of the column definition contain data that is provided by the WebLogic Server MBean interface. Refer to WebLogic documentation for more information regarding these columns.

### Fields and Data

This display includes:

- Connection** Select the connection for which you want to view collection data.
- Domain** Select the domain containing the WebLogic Server for which you want to view data.

- WLS Server** Select the WebLogic server containing the JMS Server for which you want to view data.
- JMS Server** Select the JMS Server containing the destination for which you want to view data
- Type** Select the type of destination (queue or topic) for which you want to view data, or select **All Types**.
- Module** Select the module containing the destination for which you want to view data.

### JMS Destinations on the selected JMS Server Table

This table describes all destinations on the selected JMS Server.

<b>Destination Name</b>	The name of the destination.
<b>Type</b>	The type of destination (queue or topic).
<b>Bytes Current</b>	The current number of bytes stored in the destination.*
<b>Bytes High</b>	The highest number of bytes stored in the destination since the last polling update.*
<b>Bytes Pending</b>	The current number of pending bytes stored in the destination.*
<b>Bytes Rcvd</b>	The number of bytes received in this destination since the last polling update.*
<b>Rcvd Delta</b>	The increase in the amount of bytes received (from the previous polling period to the current polling period).
<b>Rcvd Rate</b>	The rate of bytes received (per second) into the destination.*
<b>Consumers Current</b>	The current number of consumers accessing the destination.*
<b>Consumers High</b>	The highest number of consumers accessing the destination since the last polling update.*
<b>Consumers Total</b>	The total number of consumers accessing the destination since the last polling update.*
<b>Msgs Current</b>	The current number of messages in the destination.*
<b>Msgs Deleted</b>	The number of messages that have been deleted from the destination.*
<b>Msgs High</b>	The highest number of messages in the destination since the last polling update.*
<b>Msgs Moved</b>	The number of moved messages in the destination since the last polling update.*
<b>Msgs Pending</b>	The current number of pending messages in the destination.*
<b>Msgs Rcvd</b>	The number of messages received in the destination since the last polling update.*
<b>Rcvd Delta</b>	The increase in the amount of messages received (from the previous polling period to the current polling period).
<b>Rcvd Rate</b>	The rate of messages received (per second) in the destination since the last polling update.*

<b>Msgs Thresh Time</b>	The amount of time in the threshold condition since the last polling update.
<b>Bytes Thresh Time</b>	The amount of time in the threshold condition since the last polling update.
<b>Paused</b>	Indicates whether or not the destination is paused at the current time.*
<b>Consumption Paused</b>	Indicates the consumption paused state of the destination.
<b>Consumption Paused State</b>	The current consumption paused state of the destination.
<b>Production Paused</b>	Indicates the production paused state of the destination.
<b>Production Paused State</b>	The current production paused state of the destination.
<b>Insertion Paused</b>	Indicates the insertion paused state of the destination.
<b>Insertion Paused state</b>	The current insertion pause state of the destination.
<b>State</b>	The current health state of the destination.
<b>JMS Server Runtime</b>	The name of the JMS Server.
<b>time_stamp</b>	The date and time this row of data was last updated.

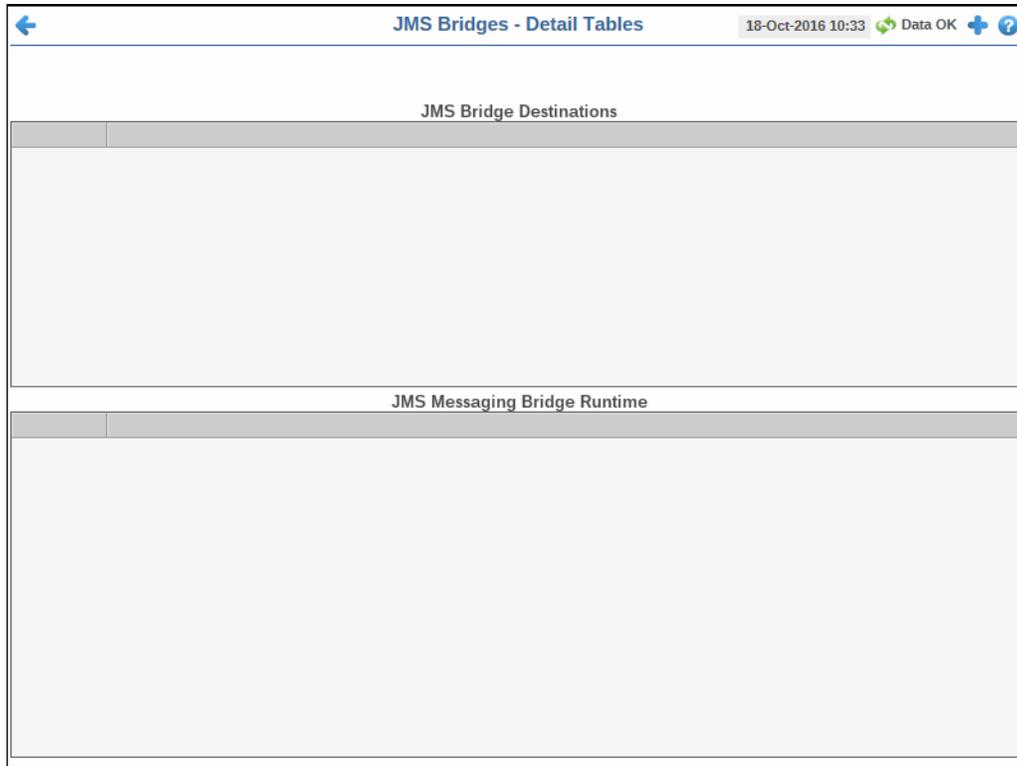
## JMS Bridges View

This view contains tabular and summary displays that show performance metrics for JMS bridges. The following displays are available:

- ["All Bridges Table"](#): A tabular view that allows you to view performance and utilization metrics for all JMS Bridge Destinations.
- ["Bridge Summary"](#): A summary view that allows you to view performance and utilization metrics for a bridge associated with a particular JMS Server.

## All Bridges Table

This table allows you to view performance and utilization metrics for all JMS Bridge Destinations.



**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- The number of items currently in the display.
- Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

**Note:** Row columns in this table with an asterisk (\*) at the end of the column definition contain data that is provided by the WebLogic Server MBean interface. Refer to WebLogic documentation for more information regarding these columns.

### Fields and Data

This display includes:

#### JMS Bridge Destinations Table

This table describes all JMS Bridge Destinations.

<b>Name</b>	The user-specified name of the instance.
<b>Source Destination Name</b>	The name of the source destination from which the messaging bridge instance reads messages.
<b>Target Destination Name</b>	The name of the target destination where the messaging bridge instance sends the messages it receives from the source destination.
<b>Async Enabled</b>	When checked, denotes that the messaging bridge instance forwards in asynchronous messaging mode.*
<b>Batch Interval</b>	The maximum amount of time, in milliseconds, that a messaging bridge instance waits before sending a batch of messages in one transaction.*
<b>Batch Size</b>	The number of messages that are processed within one transaction.*
<b>Deployment Order</b>	The priority that the server uses to determine when it deploys an item.*
<b>Durability Enabled</b>	Denotes whether or not the messaging bridge allows durable messages.*
<b>Idle Time Maximum</b>	The maximum amount of time, in seconds, that a message bridge instance remains idle.*
<b>Preserve Message Property</b>	Specifies if message properties are preserved when messages are forwarded by a bridge instance.*
<b>QOS Degradation Allowed</b>	Specifies if this messaging bridge instance allows the degradation of its quality of service (QOS) when the configured QOS is not available.*
<b>Quality Of Service</b>	The quality of service for this messaging bridge instance.*
<b>Reconnect Delay Increase</b>	The incremental delay time, in seconds, that a messaging bridge instance increases its waiting time between one failed reconnection attempt and the next retry.*
<b>Reconnect Delay Maximum</b>	The longest time, in seconds, that a messaging bridge instance waits between one failed attempt to connect to the source or target and the next retry.*
<b>Reconnect Delay Minimum</b>	The minimum amount of time, in seconds, that a messaging bridge instance waits before it tries to reconnect to the source or target destination after a failure.*
<b>Selector</b>	The filter for messages that are sent across the messaging bridge instance.*
<b>Started</b>	Specifies the initial operating state of a targeted messaging bridge instance.*
<b>time_stamp</b>	The date and time this row of data was last updated.
<b>Transaction Timeout</b>	The amount of time, in seconds, that the transaction manager waits for each transaction before timing it out.*
<b>Source Destination</b>	The source destination from which the messaging bridge instance reads messages.*
<b>Target Destination</b>	The target destination where the messaging bridge instance sends the messages it receives from the source destination.*

### JMS Messaging Bridge Runtime Table

This table describes the messaging bridge metrics.

<b>Name</b>	The name of the messaging bridge.*
<b>Location</b>	The location of the messaging bridge.*
<b>State</b>	The current state of the messaging bridge. If <b>Inactive</b> , the color of this row is set to red.*
<b>Description</b>	If the message bridge is not running, the reason why it is not running is listed here.*
<b>time_stamp</b>	The date and time this row of data was last updated.

### Bridge Summary

This summary view allows you to view performance and utilization metrics for a bridge associated with a particular JMS Server.

←
**JMS Bridge - Summary by Bridge**
18-Oct-2016 10:38 Data OK + ?

---

Domain:

---

WLS Server: 
JMS Server:

---

Bridge Name:

Source Destination:

Target Destination:

---

AsyncEnabled

Batch Interval:

Reconnect Delay Increase:

DurabilityEnabled

Batch Size:

Reconnect Delay Maximum:

PreserveMsgProperty

Idle Time Maximum:

Reconnect Delay Minimum:

QOSDegradationAllowed

Transaction Timeout:

Quality Of Service:

Started

Time Stamp:

Deployment Order:

Selector:

---

**JMS Messaging Bridge Runtime**

**Title Bar** (possible features are):

-   Open the previous and upper display.
-  Open an instance of this display in a new window.
-  Open the online help page for this display.
- open commonly accessed displays.
- The number of items currently in the display.

 **Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

 Open the **Alert Views - RTView Alerts Table** display.

---

**Note:** Row columns in this table with an asterisk (\*) at the end of the column definition contain data that is provided by the WebLogic Server MBean interface. Refer to WebLogic documentation for more information regarding these columns

---

### Fields and Data

This display includes:

<b>Domain</b>	Select the domain containing the WebLogic Server for which you want to view data.
<b>WLS Server</b>	Select the WebLogic server containing the JMS Server for which you want to view data.
<b>JMS Server</b>	Select the JMS Server containing the destination for which you want to view data
<b>Bridge Name</b>	The name of the bridge associated with the JMS Server.
<b>Source Destination</b>	The source destination from which the messaging bridge instance reads messages.*
<b>Target Destination</b>	The target destination where the messaging bridge instance sends the messages it receives from the source destination.*
<b>Async Enabled</b>	When checked, denotes that the messaging bridge instance forwards in asynchronous messaging mode.*
<b>Durability Enabled</b>	Denotes whether or not the messaging bridge allows durable messages.*
<b>Preserve Msg Property</b>	Specifies if message properties are preserved when messages are forwarded by a bridge instance.*
<b>QOS Degradation Allowed</b>	Specifies if this messaging bridge instance allows the degradation of its quality of service (QOS) when the configured QOS is not available.*
<b>Started</b>	Specifies the initial operating state of a targeted messaging bridge instance.*
<b>Batch Interval</b>	The maximum amount of time, in milliseconds, that a messaging bridge instance waits before sending a batch of messages in one transaction.*
<b>Batch Size</b>	The number of messages that are processed within one transaction.*
<b>Idle Time Maximum</b>	The maximum amount of time, in seconds, that a message bridge instance remains idle.*
<b>Transaction Timeout</b>	The amount of time, in seconds, that the transaction manager waits for each transaction before timing it out.*

<b>time_stamp</b>	The date and time the data in this display was last updated.
<b>Reconnect Delay Increase</b>	The incremental delay time, in seconds, that a messaging bridge instance increases its waiting time between one failed reconnection attempt and the next retry.*
<b>Reconnect Delay Maximum</b>	The longest time, in seconds, that a messaging bridge instance waits between one failed attempt to connect to the source or target and the next retry.*
<b>Reconnect Delay Minimum</b>	The minimum amount of time, in seconds, that a messaging bridge instance waits before it tries to reconnect to the source or target destination after a failure.*
<b>Quality of Service</b>	The quality of service for this messaging bridge instance.*
<b>Deployment Order</b>	The priority that the server uses to determine when it deploys an item.*
<b>Selector</b>	The filter for messages that are sent across the messaging bridge instance.*

#### JMS Messaging Bridge Runtime Table

This table describes the messaging bridge metrics.

<b>Location</b>	The location of the messaging bridge.*
<b>State</b>	The current state of the messaging bridge. If <b>Inactive</b> , the color of this row is set to red.*
<b>Description</b>	If the message bridge is not running, the reason why it is not running is listed here.*

## JMS Connections View

This view includes a display that presents a view of JMS performance metrics for a particular server.

- [“JMS Connections - Detail by Server”](#): A series of tables that allow you to view performance and utilization metrics for all JMS connections, JMS consumers, JMS producers, and JMS durable subscribers on a particular JMS server.

## JMS Connections - Detail by Server

Track performance and utilization metrics for all JMS connections, JMS consumers, JMS producers, and JMS durable subscribers on a particular JMS server.

**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- 6,047** The number of items currently in the display.
- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

**Note:** Row columns in this table with an asterisk (\*) at the end of the column definition contain data that is provided by the WebLogic Server MBean interface. Refer to WebLogic documentation for more information regarding these columns

### Fields and Data

This display includes:

- Domain** Select the domain containing the WebLogic Server for which you want to view data.

**WLS Server** Select the WebLogic server containing the JMS Server for which you want to view data.

**JMS Server** Select the JMS Server containing the destination for which you want to view data

### JMS Connections Table

This table describes all JMS connections on the selected JMS server.

<b>Name</b>	The name of the connection.
<b>ClientID</b>	The client id for the connection.*
<b>HostAddress</b>	The host address of the client JVM as a string.*
<b>Sessions Current Count</b>	The current number of sessions for the connection.*
<b>Sessions High Count</b>	The highest number of sessions for the connection since the last polling update.
<b>Sessions Total Count</b>	The total number of sessions for the connection since the last polling update.
<b>time_stamp</b>	The date and time this row of data was last updated.

### JMS Consumers Table

This table describes all JMS consumers on the selected JMS server.

<b>Name</b>	The name of the consumer.
<b>Active</b>	When checked, denotes that the consumer is active.*
<b>Durable</b>	When checked, denotes that the consumer is durable.*
<b>Bytes Pending Count</b>	The number of bytes pending by the consumer.*
<b>Bytes Received Count</b>	The number of bytes received by the consumer since the last polling update.*
<b>Delta</b>	The increase in the amount of bytes received (from the previous polling period to the current polling period).
<b>Rate</b>	The rate of bytes received (per second) by the consumer.
<b>Messages Pending Count</b>	The number of pending messages by the consumer.*
<b>Messages Received Count</b>	The number of messages received by the consumer since the last polling update.*
<b>Delta</b>	The increase in the amount of messages received (from the previous polling period to the current polling period).
<b>Rate</b>	The number of messages received per second by the consumer.
<b>JMS Connection Runtime</b>	The name of the JMS connection.
<b>JMS Session Runtime</b>	The name of the JMS session.
<b>Destination Name</b>	The name of the destination for the consumer.

<b>Selector</b>	The name of the selector associated with the consumer.
<b>time_stamp</b>	The date and time this row of data was last updated.

**JMS Producers Table**

This table describes all JMS producers on the selected JMS server.

<b>Name</b>	The name of the producer.
<b>Bytes Pending Count</b>	The current number of bytes that are pending by the producer.*
<b>Bytes Sent Count</b>	The number of bytes sent by the producer since the last polling update.*
<b>Delta Bytes Sent Count</b>	The increase in the amount of bytes sent (from the previous polling period to the current polling period).
<b>Rate Bytes Sent Count</b>	The number of bytes sent per second by the producer.
<b>Messages Pending Count</b>	The number of messages pending by the producer.*
<b>Messages Sent Count</b>	The number of messages sent by the producer since the last polling update.*
<b>Delta Messages Sent Count</b>	The increase in the amount of messages sent (from the previous polling period to the current polling period).
<b>Rate Messages Sent Count</b>	The number of messages sent by the producer per second.
<b>JMS Connection Runtime</b>	The name of the JMS connection.
<b>JMS Session Runtime</b>	The name of the JMS session.
<b>time_stamp</b>	The date and time this row of data was last updated.

**JMS Durable Subscribers Table**

This table describes all JMS durable subscribers on the selected JMS server.

<b>ClientID</b>	The client ID for the durable subscriber.*
<b>Active</b>	When checked, denotes that the subscription is being used by a durable subscriber.*
<b>No Local</b>	Indicates whether the durable subscriber receives local messages that it has published.*
<b>Bytes Current Count</b>	The number of bytes received by the durable subscriber.*
<b>Bytes Pending Count</b>	The current number of bytes that are pending by the durable subscriber.*
<b>Messages Current Count</b>	The number of messages still available by this durable subscriber.*
<b>Messages Deleted Current Count</b>	The number of messages deleted from the destination.*
<b>Messages High Count</b>	The highest number of messages for the durable subscriber since the last polling update.

<b>Messages Moved Current Count</b>	The number of messages that have been moved from the destination.*
<b>Messages Pending Count</b>	The number of messages pending by this durable subscriber.*
<b>Messages Received Count</b>	The number of messages received by the durable subscriber since the last polling update.
<b>Current Consumer Info Client ID</b>	The client ID for the consumer.*
<b>Current Consumer Info Connection Address</b>	The connection address of the consumer.*
<b>Current Consumer Info Durable</b>	When checked, denotes that the current consumer is durable.*
<b>Current Consumer Info Name</b>	The name of the current consumer.*
<b>Current Consumer Info No Local</b>	Indicates whether the consumer receives local messages it has published itself.*
<b>Current Consumer Info Selector</b>	The message selector defined for this consumer.*
<b>Name</b>	The name of the durable subscriber.*
<b>Subscription Name</b>	The subscription name for the durable subscriber.*
<b>Selector</b>	The message selector defined for the durable subscriber.*
<b>time_stamp</b>	The date and time this row of data was last updated.

---

## Oracle WebLogic - HTML

The HTML version features an overview display, "[WebLogic Overview Display - HTML](#)" (pictured below), and the following Views which can be found under **Components** tab > **Application/ Web Servers > Oracle WebLogic**:

- ["WebLogic Servers View - HTML"](#)
- ["Single WebLogic Server View - HTML"](#)
- ["WebLogic Applications View - HTML"](#)

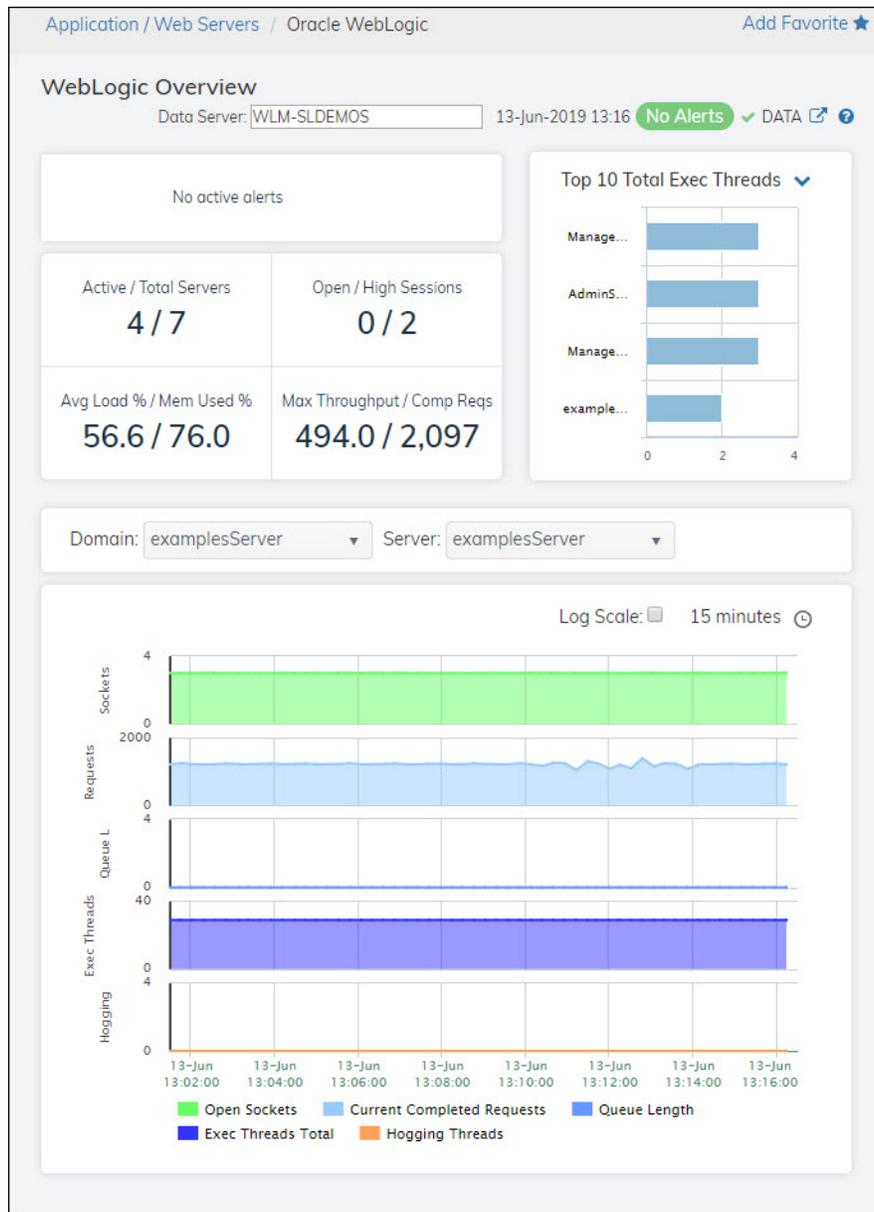
## WebLogic Overview Display - HTML

The **WebLogic Overview** is the top-level display for the Oracle WebLogic Solution Package, which provides a good starting point for immediately getting the status of all your servers on your Data Server. You can select the RTView DataServer for which you want to see data and easily view the current data for that DataServer including:

- The total number of active alerts for the selected DataServer, including the total number of critical and warning alerts.
- The number of active and total servers.
- The current number of open sessions and the highest number of registered sessions on your connected DataServer.
- The maximum average load percentage and maximum memory used percentage across all servers on your connected DataServer.
- The maximum throughput and maximum number of completed requests across all threads on your connected DataServer.
- A visual list of the top 10 servers with the highest number of execution threads and the most current pending requests on your connected DataServer.

You can hover over each region in the upper half of the Overview to see more detail. You can also drill down to see even more detail by clicking on each respective region in the Overview. For example, clicking on the alerts in the **CRITICAL** and **WARNING** alerts region opens the **Alerts Table by Components** display.

The bottom half of the display provides a trend graph for a particular server. You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.



## WebLogic Servers View - HTML

These displays present performance metrics and alert statuses for all Oracle WebLogic Servers and Clusters. Clicking **WebLogic Servers** from the left menu opens the ["WebLogic Servers Table - HTML"](#) display, which shows all available utilization metrics for all Oracle® WebLogic servers in a tabular format. The following displays are available:

- **Servers Heatmap:** Opens the ["WebLogic Servers Heatmap - HTML"](#) display, which shows status and alerts for all Oracle® WebLogic servers in a heatmap.
- **Clusters Table:** Opens the ["WebLogic Clusters Table - HTML"](#) display, which enables you to track utilization and performance metrics for all clusters on a particular domain, or on all domains.

## WebLogic Servers Table - HTML

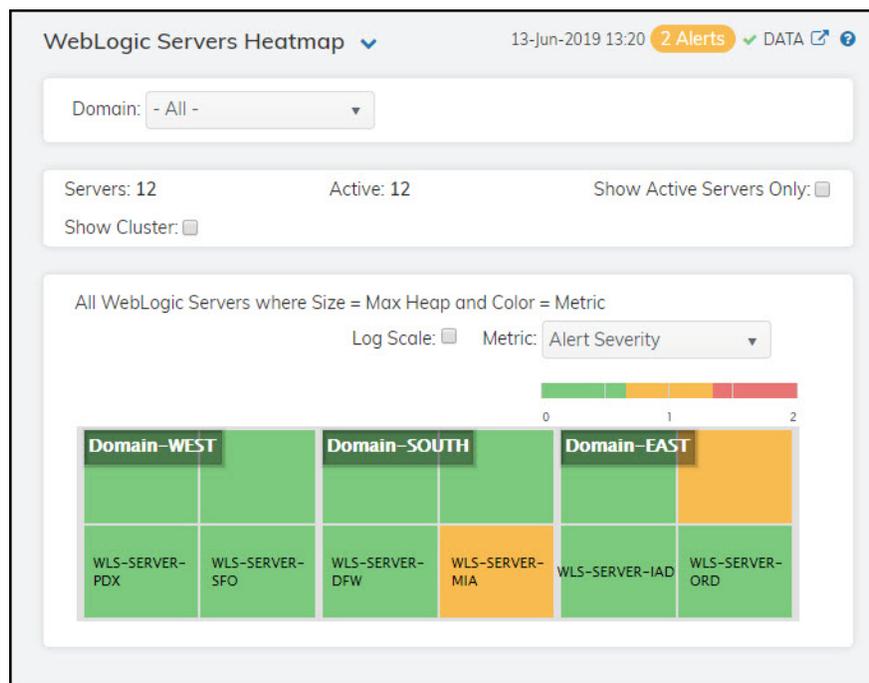
This display provides utilization metrics for all WebLogic Servers for a particular domain in a tabular format. Each row in this table includes heap, processing, thread, and version metrics (among others) for a particular server. Click a column header to sort column data in numerical or alphabetical order. Double-click on a table row to drill-down to the ["WebLogic Server Summary - HTML"](#) display and view metrics for that particular server. You can click on the drop down list in the display title to toggle between commonly accessed displays.

Domain	Cluster	Server	State
Domain-WEST	Cluster-LAX	WLS-SERVER-LAX	RUNNING
Domain-WEST	Cluster-PDX	WLS-SERVER-PDX	RUNNING
Domain-WEST	Cluster-SEA	WLS-SERVER-SEA	RUNNING
Domain-WEST	Cluster-SFO	WLS-SERVER-SFO	RUNNING
Domain-SOUTH	Cluster-ATL	WLS-SERVER-ATL	RUNNING
Domain-SOUTH	Cluster-DFW	WLS-SERVER-DFW	RUNNING

## WebLogic Servers Heatmap - HTML

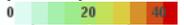
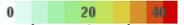
Clicking **Servers Heatmap** in the left/navigation menu opens the **WebLogic Servers Heatmap** display, which allows you to view the status and alerts of all Oracle® WebLogic servers. You can view the servers in the heatmap based on the following metrics: Alert Severity, Alert Count, Jvm CPU %, Host CPU %, Jvm Memory %, Open Sockets, Thread Total Count, and Hogging Threads.

The heatmap is organized by host, each rectangle representing a server. The rectangle color indicates the most critical alert state. Click on a node to drill-down to the [“WebLogic Server Applications Summary - HTML”](#) display and view metrics for a particular connection. You can toggle between the commonly accessed displays by clicking the display title’s drop down list. Mouse-over rectangles to view more details about host performance and status.



### Available Metrics

Select the metric driving the heatmap display. The default is Alert Severity. Each **Metric** has a color gradient bar that maps values to colors. The heatmap organizes the instances by host, where each rectangle represents an instance. Mouse-over any rectangle to display the current values of the metrics for the instance. Click on a rectangle to drill-down to the associated [“WebLogic Server Applications Summary - HTML”](#) display for a detailed view of metrics for that particular domain.

<b>Alert Severity</b>	<p>The maximum alert level in the item (index) associated with the rectangle. Values range from <b>0</b> to <b>2</b>, as indicated in the color gradient bar , where <b>2</b> is the greatest <b>Alert Severity</b>.</p> <p><b>2</b> -- Metrics that have exceeded their specified <b>ALARMLEVEL</b> threshold and have an Alert Severity value of <b>2</b> are shown in red. For a given rectangle, this indicates that one or more metrics have exceeded their alarm threshold.</p> <p><b>1</b> -- Metrics that have exceeded their specified <b>WARNINGLEVEL</b> threshold and have an Alert Severity value of <b>1</b> are shown in yellow. For a given rectangle, this indicates that one or more metrics have exceeded their warning threshold.</p> <p><b>0</b> -- Metrics that have not exceeded either specified threshold have an Alert Severity value of <b>0</b> and are shown in green. For a given rectangle, this indicates that no metrics have exceeded a specified alert threshold.</p>
<b>Alert Count</b>	<p>The total number of alarm and warning alerts in a given item (index) associated with the rectangle.</p> <p>The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>Jvm CPU %</b>	<p>The percentage of JVM CPU currently being used in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the alert threshold of <b>WisServerCpuHigh</b>. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>Host CPU %</b>	<p>The percentage of Host CPU currently being used in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the alert threshold of <b>WisServerHostCpuHigh</b>. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>JVM Memory %</b>	<p>The percentage of JVM Memory currently being used in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the alert threshold of <b>WisMemoryUsageHigh</b>. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>Open Sockets</b>	<p>The total number of open sockets currently being used in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the alert threshold of <b>WisOpenSocketsHigh</b>. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>Thread Total Count</b>	<p>The total number of threads in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the alert threshold of <b>WisThreadsTotalHigh</b>. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>Hogging Threads</b>	<p>The total number of hogging threads currently being used in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the alert threshold of <b>WisHoggingThreadsHigh</b>. The middle value in the gradient bar indicates the middle value of the range.</p>

## WebLogic Clusters Table - HTML

Clicking **Clusters Table** in the left/navigation menu opens the **WebLogic Clusters Table** display, which allows you to track utilization and performance metrics for all clusters on a particular domain, or on all domains. Each row in this table includes metrics for a particular cluster. Click a column header to sort column data in numerical or alphabetical order. Double-click on a table row to drill-down to the ["WebLogic Servers Table - HTML"](#) display and view metrics for that particular server. You can click on the drop down list in the display title to toggle between commonly accessed displays.

Domain	Cluster	Server Count	Running Count	Name
TestDomain		1	1	
TestDomain	Cluster-114-2	2	0	
TestDomain	TestCluster-114	3	2	
examplesServer		1	1	

## Single WebLogic Server View - HTML

These displays present detailed performance metrics and alert statuses for a particular WebLogic server. Clicking **Single WebLogic Server** from the left menu opens the ["WebLogic Server Summary - HTML"](#) display, which allows you to track utilization, performance, and trend data for a particular WebLogic server. The following displays are available in this View:

- **JVM Summary:** Opens the ["WebLogic JVM Summary - HTML"](#) display, which displays the JVM details for a particular WebLogic server on a specific domain.
- **Server Detail:** Opens the ["WebLogic Server Detail Tables - HTML"](#) display, which displays server runtime data, threadpool runtime data, JRockit runtime data, and server version information for a specific WebLogic server
- **JDBC Summary:** Opens the ["WebLogic JDBC Summary - HTML"](#) display, which displays JDBC module utilization, performance, and trend data for a specific WebLogic server.
- **Work Manager:** Opens the ["WebLogic Server Work Manager Table - HTML"](#) display, which displays server runtime data for all work managers on a specific WebLogic Server.
- **Persistent Stores:** Opens the ["WebLogic JMS Persistent Stores Detail Tables - HTML"](#) display, which displays available utilization and performance data for all configurations on a specific domain.

## WebLogic Server Summary - HTML

Clicking **Single WebLogic Server** in the left/navigation menu opens the **WebLogic Server Summary** display, which allows you to track utilization, performance, and trend data for a particular WebLogic server. Hovering over the information boxes at the top of the display provides additional details, and clicking on them takes you to the ["WebLogic Servers Table - HTML"](#) display, where you can view additional data.

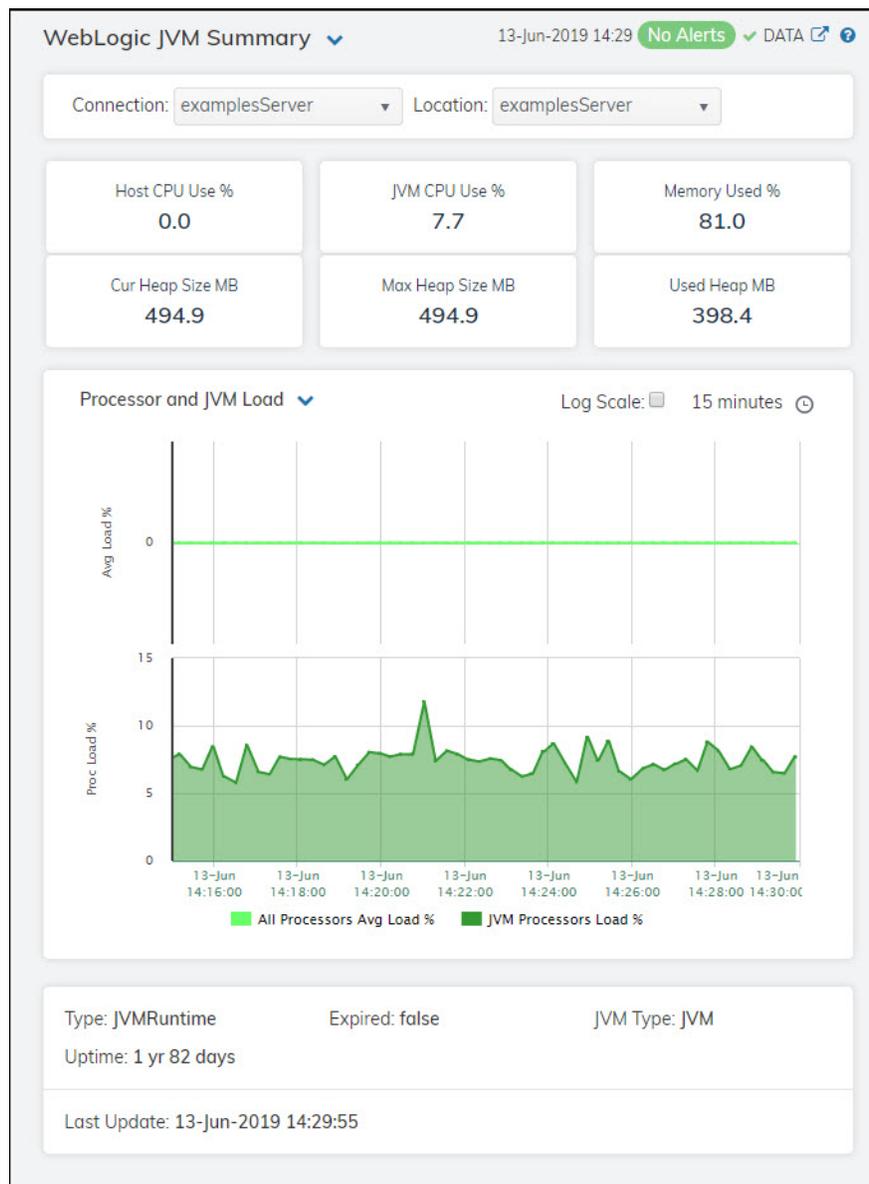
In the trend graph region, you can view open socket, completed requests, queue length, execute threads, and hogging threads trend data for the selected connection/location combination.



## WebLogic JVM Summary - HTML

Clicking **JVM Summary** from the left/navigation menu opens the **WebLogic JVM Summary** display, which allows you to view the JVM details for a particular WebLogic server on a specific domain. Hovering over the information boxes at the top of the display provides additional details, and clicking on them takes you to the "[WebLogic Servers Table - HTML](#)" display, where you can view additional data.

There are two options in the trend graph region: **Processor and JVM Load** and **Heap Memory**. Selecting **Processor and JVM Load** traces the percentage of the host CPU being used and the percentage of the JVM CPU being used for the selected connection/location combination over a defined time range. Selecting **Heap Memory** traces the maximum amount of heap available, the current size of the JVM heap, and the total amount of used heap memory for the selected connection/location combination over a defined time range.



## WebLogic Server Detail Tables - HTML

Clicking **Server Detail** in the left/navigation menu opens the **WebLogic Server Detail Tables** display, which allows you to view server runtime, threadpool runtime, JRockit runtime, and server version information for a specific WebLogic server. Double-clicking on a row in any of the tables opens the "[WebLogic Server Summary - HTML](#)" display, which allows you to view additional details for the server.

WebLogic Server Detail Tables 13-Jun-2019 14:30 No Alerts DATA [?](#)

Connection: TestDomain Location: AdminServer

---

**Server Runtime**

Domain	Server	Cluster	Activation Time	Av
TestDomain	AdminServer		49 years	

---

**Threadpool Runtime**

Connection	Server	Completed Requests	Current Completed Reque...	Exec Idle Threads	Exec Thr Tot
TestDomain	AdminServer	1,134,632	2,095	2	

---

**JRockit Runtime**

Connection	Server	All Proc Avg Load	Concurrent	FreeHeap	FreePhy
TestDomain	AdminServer	30.75		43525360	

---

**Server Version Info**

WebLogic Server 10.3.3.0 Fri Apr 9 00:05:28 PDT 2010 1321401

## WebLogic JDBC Summary - HTML

Clicking **JDBC Summary** in the left/navigation menu opens the **WebLogic JDBC Summary** display, which provides a view of JDBC module utilization, performance data, and trend data for a specific WebLogic server. Hovering over the information boxes at the top of the display provides additional details, and clicking on them takes you to the "[WebLogic Servers Table - HTML](#)" display, where you can view additional data.

There are two options in the trend graph region: **Connections** and **Rates**. Selecting **Connections** traces the highest number of active database connections in the instance of the data source since the data source was instantiated, the number of connections currently in use by applications, and the number of prepared and callable statements currently cached in the statement cache for the selected module over a defined time range. Selecting **Rates** traces the cumulative, running count of the number of times that the statement cache was accessed, the number of times (per second) that a statement request could not be satisfied with a statement from the cache, the cumulative (running) count of the number of statements added to the statement cache, and the cumulative (running) count of statements discarded from the cache for the selected module over a defined time range.

WebLogic JDBC Summary ▼
13-Jun-2019 14:31 No Alerts DATA [🔗](#) [?](#)

Domain: examplesServer ▼
Server: examplesServer ▼

Module: examples-demo ▼

Active Connections  
**2**

Reconnect Failures  
**0**

Hogging Threads  
**1**

Thread Queue Length  
**3**

Execute Threads  
**0**

Waiting for Connection Current Count  
**0**

Connections ▼
Log Scale:  15 minutes 🕒

High Active Connections
  Current Active Connections
  Current Size

State: Running	Enabled: true	Active Conns Current: 2
Active Conns High: 3	Connections Total Count: 3	Current Capacity: 2
Failures to Reconnect: 0	Leaked Conns: 0	Stmt Cache Current Size: 6
Stmt Cache Hits: 0	Stmt Cache Missed: 6	Stmt Cache Accesses: 6
Stmt Cache Additions: 6	Stmt Cache Deletions: 0	Reserve Request: 10
Failed Reserve Requests: 0	Wait Secs High: 0	Waiting Conn: 0
Waiting Conn Fail: 0	Waiting Conn High: 0	Waiting Conn Success: 2
Waiting Conn Total: 0	Connection Delay Time: 5172	
Driver Version: org.apache.derby.jdbc.ClientDriver		

Last Update: 13-Jun-2019 14:30:58

## WebLogic Server Work Manager Table - HTML

Clicking **Work Manager** in the left/navigation menu opens the **WebLogic Server Work Manager Table** display, which allows you to view server runtime data for all work managers on a specific WebLogic Server. Double-clicking on a row in any of the tables opens the [“WebLogic Server Summary - HTML”](#) display, which allows you to view additional details for the server.

WebLogic Server Work Manager Table 13-Jun-2019 14:31 No Alerts DATA

Domain: examplesServer Server: examplesServer

Name	Application Name	Completed Requests	Cur Comp. Request
default	examples-multiDataSource-demo	0	
weblogic.wsee.mdb.DispatchPolicy	asyncServletEar	0	
default	esphere	0	
weblogic.logging.DomainLogBroad		88	
default	examples-demo	0	
default	jdbcRowSetsEar	0	
weblogic.jms.ReliableWseeSAFAge		0	
default	SampleSearchWebApp	0	

Page 1 of 2 1 - 40 of 72 items

## WebLogic JMS Persistent Stores Detail Tables - HTML

Clicking **Persistent Stores** in the left/navigation menu opens the **WebLogic JMS Persistent Stores Detail Tables** display, which allows you to view available utilization and performance data for all configurations on a specific domain.

WebLogic JMS Persistent Stores Detail Tables 13-Jun-2019 14:32 No Alerts [DATA](#) [?](#)

Domain: examplesServer

**Persistent Store Runtime**

Name	Server	CreateCount	DeleteC
WseeFileStore	examplesServer	0	

**Persistent Store Connection Runtime**

Name	Server	PersistentS
weblogic.messaging.examplesJMSServer.header	examplesServer	exampleJDBCStore

## WebLogic Applications View - HTML

These displays present several views of performance metrics for applications on clusters and a particular WebLogic server. Clicking **WebLogic Applications** from the left menu opens the [“WebLogic Cluster Applications Table - HTML”](#) display, which enables you to track utilization and performance metrics for all clusters on a particular domain, or on all domains.

- **Cluster App Summary:** Opens the [“WebLogic Clustered Application Summary - HTML”](#) display, which allows you to view performance and utilization metrics for all applications on a particular cluster, or for all applications on all clusters.
- **All Apps Heatmap:** Opens the [“WebLogic Server Applications Heatmap - HTML”](#) display, which shows a heatmap view of the status and alerts of all applications within a specific WebLogic server.
- **All Apps Summary:** Opens the [“WebLogic Server Applications Summary - HTML”](#) display, which allows you to track performance, utilization, and trend data for all applications on a single WebLogic server.
- **Application Summary:** Opens the [“WebLogic Application Summary - HTML”](#) display, which allows you to view performance, utilization, and trend data for a single application component on a single WebLogic server.
- **Application Trends:** Opens the [“WebLogic Application Metric Trends - HTML”](#) display, which allows you to view trend data for a single application on a particular WebLogic server.
- **App Components Heatmap:** Opens the [“WebLogic Application Components Heatmap - HTML”](#) display, which provides a heatmap view of the status and alerts of all application components contained within each application on a particular WebLogic server.
- **App Component Summary:** Opens the [“WebLogic Application Component Summary - HTML”](#) display, which allows you to view performance, utilization, and trend data for a particular application component on a WebLogic Server.

## WebLogic Cluster Applications Table - HTML

Clicking **WebLogic Applications** in the left/navigation menu opens the **WebLogic Cluster Applications Table** display, which allows you to view performance and utilization metrics for all applications on a particular cluster, or for all applications on all clusters. Double-clicking on a row in the table opens the "[WebLogic Clustered Application Summary - HTML](#)" display, which allows you to view additional information for that particular application.

WebLogic Cluster Applications Table 13-Jun-2019 14:32 DATA

Domain: examplesServer Cluster: - All -

Clusters: 1

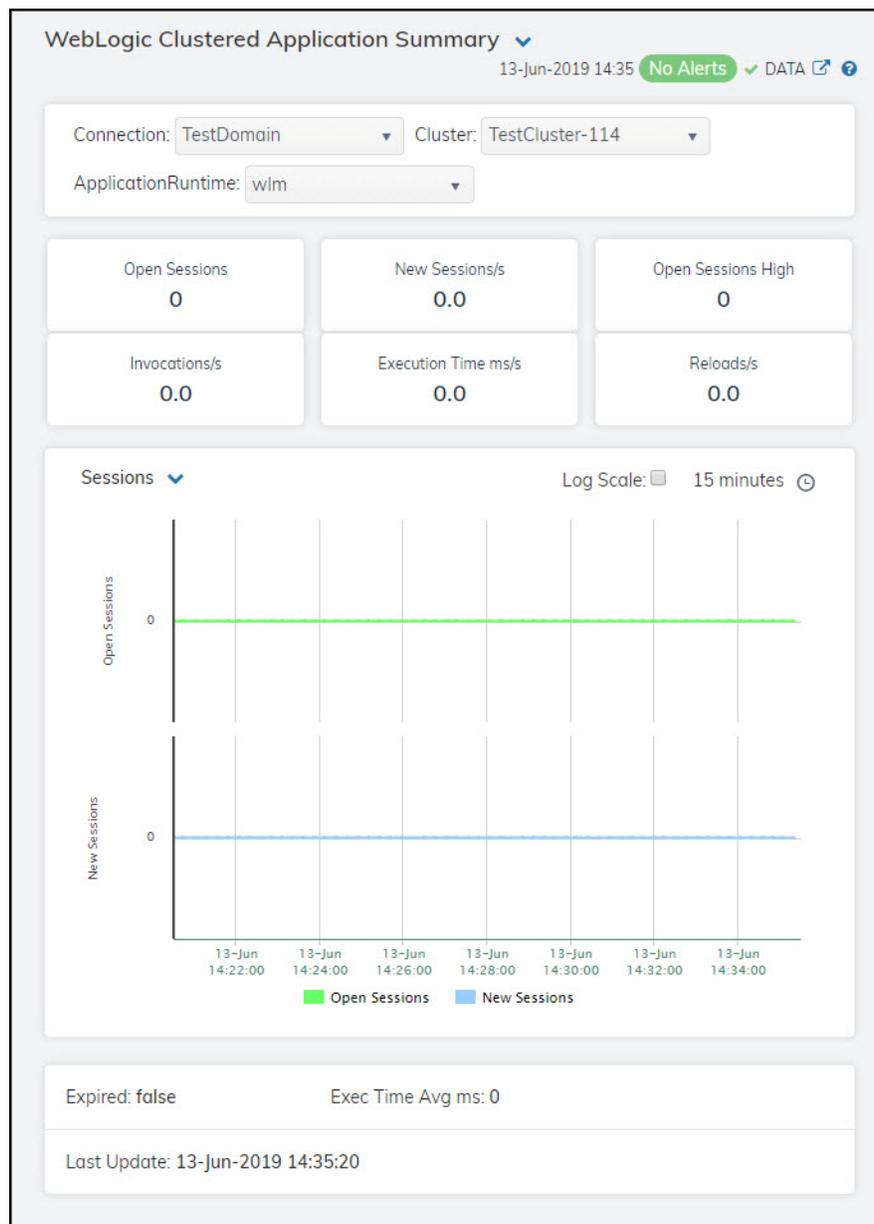
Domain	Cluster	Application	Open Sessions
examplesServer		xmlBeanEar	

Page 1 of 8 1 - 40 of 289 items

## WebLogic Clustered Application Summary - HTML

Clicking **Cluster App Summary** in the left/navigation menu opens the **WebLogic Clustered Application Summary** display, which allows you to view session information for a particular application. Hovering over the information boxes at the top of the display provides additional details, and clicking on them takes you to the "[WebLogic Servers Table - HTML](#)" display, where you can view additional data.

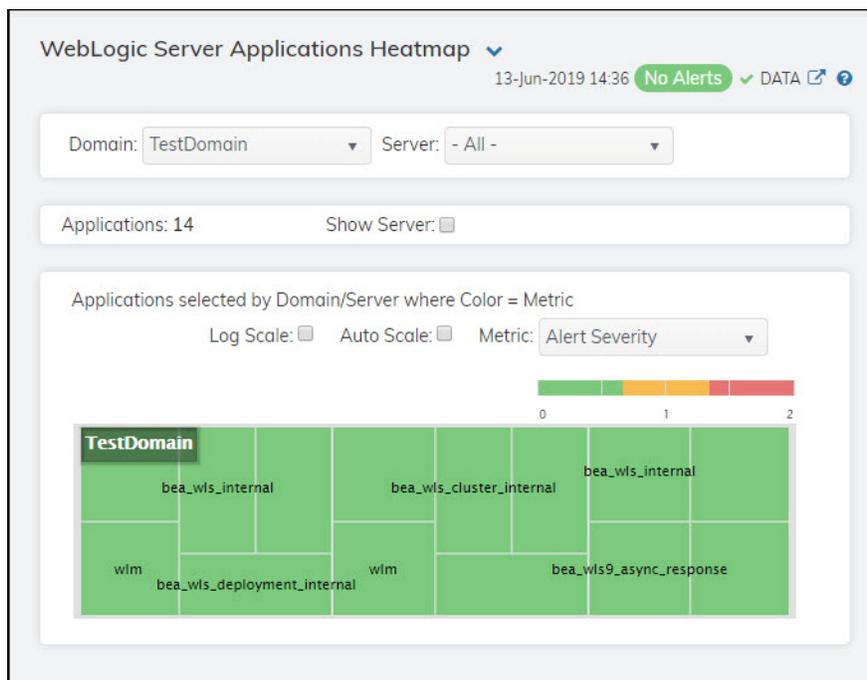
There are two options in the trend graph region: **Sessions** and **Rates**. Selecting **Sessions** traces the number of open sessions and the number of newly created sessions over a defined time range. Selecting **Rates** traces the number of sessions created per second, the number of invocations per second, and the execution time in milliseconds per second over a defined time range.



## WebLogic Server Applications Heatmap - HTML

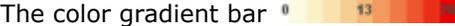
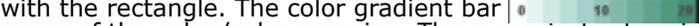
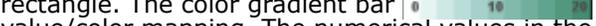
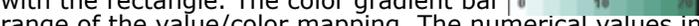
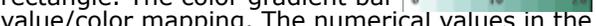
Clicking **All Apps Heatmap** in the left/navigation menu opens the **WebLogic Servers Applications Heatmap** display, which provides a heatmap view of the status and alerts of all applications within a specific WebLogic server. You can view the servers in the heatmap based on the following metrics: alert severity, alert count, total open sessions, open sessions rate, total invocations, invocations rate, execution time, or total reloads.

The heatmap is organized by domain, where each rectangle represents an application. The rectangle color indicates the most critical alert state. Click on a node to drill-down to the ["WebLogic Server Applications Summary - HTML"](#) display and view metrics for all applications associated with that server. You can toggle between the commonly accessed displays by clicking the display title's drop down list. Mouse-over rectangles to view more details about connection performance and status.



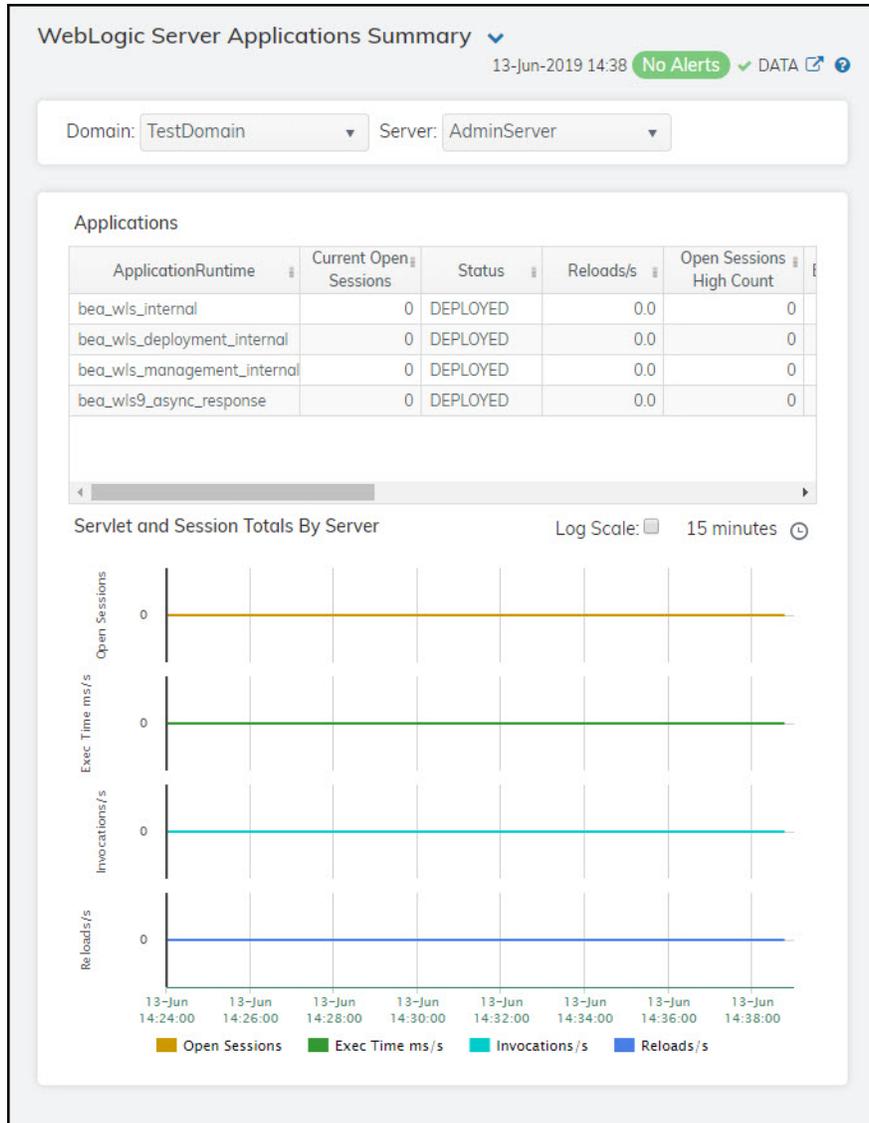
### Available Metrics

Select the metric driving the heatmap display. The default is Alert Severity. Each **Metric** has a color gradient bar that maps values to colors. The heatmap organizes the instances by domain, where each rectangle represents an application. Mouse-over any rectangle to display the current values of the metrics for the application. Click on a rectangle to drill-down to the associated ["WebLogic Server Applications Summary - HTML"](#) display for a detailed view of metrics for that particular server.

<b>Alert Severity</b>	<p>The maximum alert level in the item (index) associated with the rectangle. Values range from <b>0</b> to <b>2</b>, as indicated in the color gradient bar , where <b>2</b> is the greatest <b>Alert Severity</b>.</p> <p><b>2</b> -- Metrics that have exceeded their specified <b>ALARMLEVEL</b> threshold and have an Alert Severity value of <b>2</b> are shown in red. For a given rectangle, this indicates that one or more metrics have exceeded their alarm threshold.</p> <p><b>1</b> -- Metrics that have exceeded their specified <b>WARNINGLEVEL</b> threshold and have an Alert Severity value of <b>1</b> are shown in yellow. For a given rectangle, this indicates that one or more metrics have exceeded their warning threshold.</p> <p><b>0</b> -- Metrics that have not exceeded either specified threshold have an Alert Severity value of <b>0</b> and are shown in green. For a given rectangle, this indicates that no metrics have exceeded a specified alert threshold.</p>
<b>Alert Count</b>	<p>The total number of alarm and warning alerts in a given item (index) associated with the rectangle.</p> <p>The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>Open Sessions</b>	<p>The total number of open sessions in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from <b>0</b> to the alert threshold of <b>WisOpenSessionsHigh</b>, which is <b>10</b>. The middle value in the gradient bar indicates the middle value of the range (the default is <b>5</b>).</p>
<b>Open Sessions/s</b>	<p>The number of sessions opened per second in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of objects in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>Invocations/s</b>	<p>The number of invocations per second in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of objects in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>Invocations</b>	<p>The total number of invocations in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of objects in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>Execution Time ms</b>	<p>The execution time, in milliseconds, in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of objects in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>Reloads</b>	<p>The total reload count in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of objects in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>

## WebLogic Server Applications Summary - HTML

Clicking **All Apps Summary** in the left/navigation menu opens the **WebLogic Server Applications Summary** display, which allows you to track performance, utilization, and trend data for all applications on a single WebLogic server. This display is broken up into a table listing all applications on the selected server and a **Servlet and Session Totals By Server** trend graph that traces open sessions, execution time, invocation creation rate, and reloads creation rate for the selected server over a defined time range. Clicking a row in the table opens the ["WebLogic Application Summary - HTML"](#) display, where you can view additional data for the selected application.



## WebLogic Application Summary - HTML

Clicking **Application Summary** in the left/navigation menu opens the **WebLogic Application Summary** display, which allows you to view performance, utilization, and trend data for a single application on a single WebLogic server. Hovering over the information boxes at the top of the display provides additional details.

The **Application Servlets** table lists all associated servlets and their various metrics. Double-clicking on a row in this table opens the ["WebLogic Server Applications Summary - HTML"](#) display where you can view additional details.

The **Application Activity Trends** trend graph traces the current number of open sessions, the execution time rate (in milliseconds/second), the invocation creation rate, and the reloads rate over a defined time range.

WebLogic Application Summary ▼
13-Jun-2019 14:40 No Alerts ✓ DATA [🔗](#) [🔔](#)

Domain: examplesServer ▼
Server: examplesServer ▼

Application: consoleapp ▼

Open Sessions  
**0**

Session Timeout s  
**7,200**

Current Completed Requests  
**0**

Pending Requests  
**0**

Stuck Threads  
**0**

Exec Time Avg ms  
**6220**

#### Application Servlets

Servlet	Alert Level	Alert Count	Completed Requests	Cur Completed Requests	Comy Requ
FileDefault	✓		0	0	
AppManagerServlet	✓		0	0	
JspServlet	✓		0	0	
JSPCServlet	✓		0	0	
WebServiceServlet	✓		0	0	
FileServlet	✓		0	0	

#### Application Activity Trends Log Scale: 15 minutes 🕒

13-Jun 14:26:00 13-Jun 14:28:00 13-Jun 14:30:00 13-Jun 14:32:00 13-Jun 14:34:00 13-Jun 14:36:00 13-Jun 14:38:00 13-Jun 14:40:00

■ Current Open Sessions
 ■ Exec Time ms/s
 ■ Invocations/s
 ■ Reloads/s

Servlets: 11
Open Sessions High Count: 1
Invocations: 9

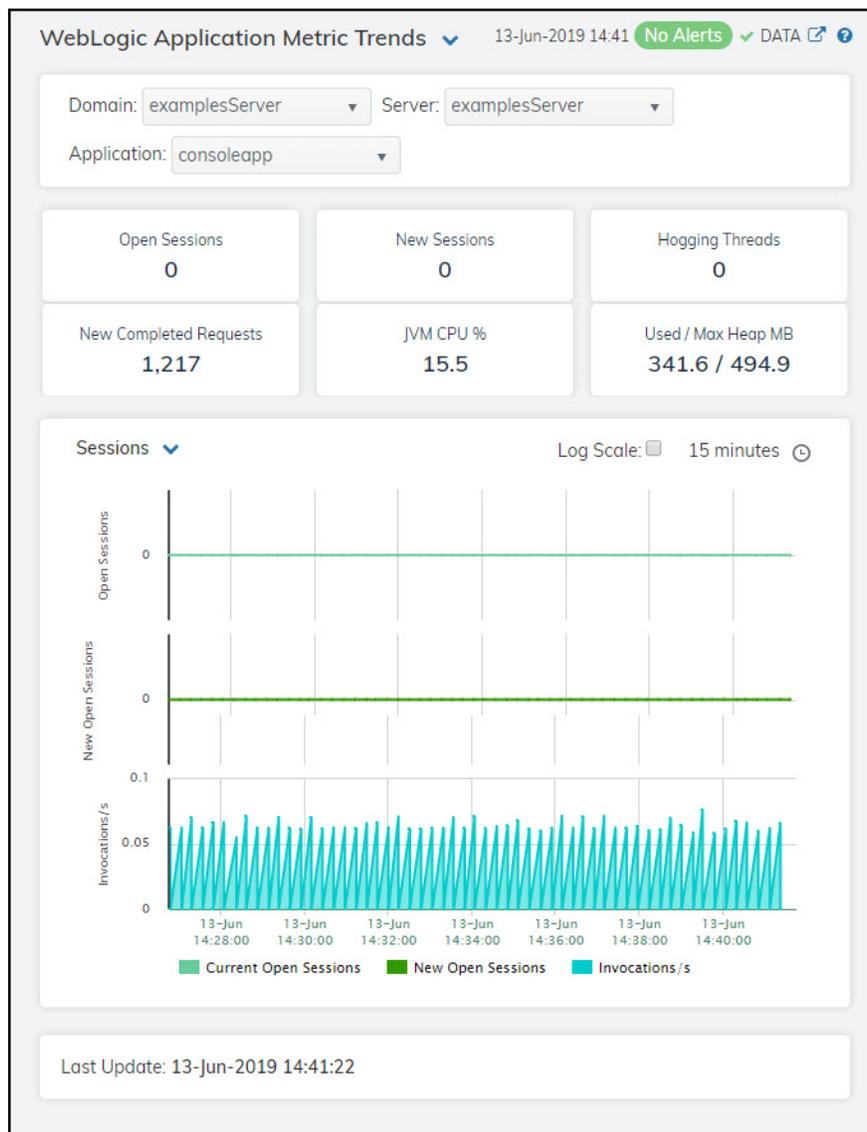
Reloads: 10
Exec Time Avg ms: 6220
Expired: false

Last Update: 13-Jun-2019 14:40:50

## WebLogic Application Metric Trends - HTML

Clicking **Application Trends** in the left/navigation menu opens the **WebLogic Application Metric Trends** display, which allows you to view metrics and trend data for a single application on a particular WebLogic server. Hovering over the information boxes at the top of the display provides additional details.

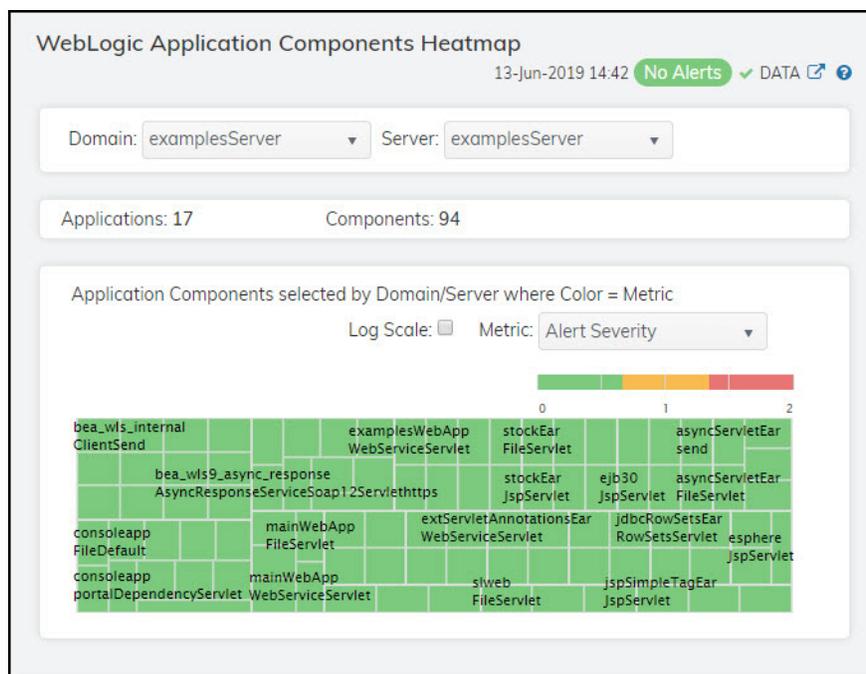
The trend graph at the bottom of the display has three options: **Sessions**, **Threads**, and **JVM**. **Sessions** traces the current number of open sessions, the number of new open sessions, and the invocation rate over a defined time range. **Threads** traces the current number of thread requests and the number of hogging threads over a defined time range. **JVM** traces the JVM usage percentage, the maximum heap, and the used heap over a defined time range.



## WebLogic Application Components Heatmap - HTML

Clicking **Apps Components Heatmap** in the left/navigation menu opens the **WebLogic Applications Components Heatmap** display, which provides a heatmap view of the status and alerts of all application components contained within each application on a particular WebLogic server. You can view the applications in the heatmap based on the following metrics: alert severity, alert count, completed requests, completed requests rate, invocations rate, invocations, execution time, and reloads.

The heatmap is organized by domain/server, where each rectangle represents an application. The rectangle color indicates the most critical alert state. Click on a node to drill-down to the ["WebLogic Application Component Summary - HTML"](#) display and view metrics for the selected component. You can toggle between the commonly accessed displays by clicking the display title's drop down list. Mouse-over rectangles to view more details about component performance and status.

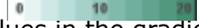
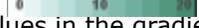


### Available Attributes

Select the metric driving the heatmap display. The default is **Alert Severity**. Each **Metric** has a color gradient bar that maps values to colors. The heatmap organizes the collections by domain/server, where each rectangle represents an application. Mouse-over any rectangle to display the current values of the metrics for the application. Click on a rectangle to drill-down to the associated ["WebLogic Application Component Summary - HTML"](#) display for a detailed view of metrics for that particular application.

#### Completed Requests

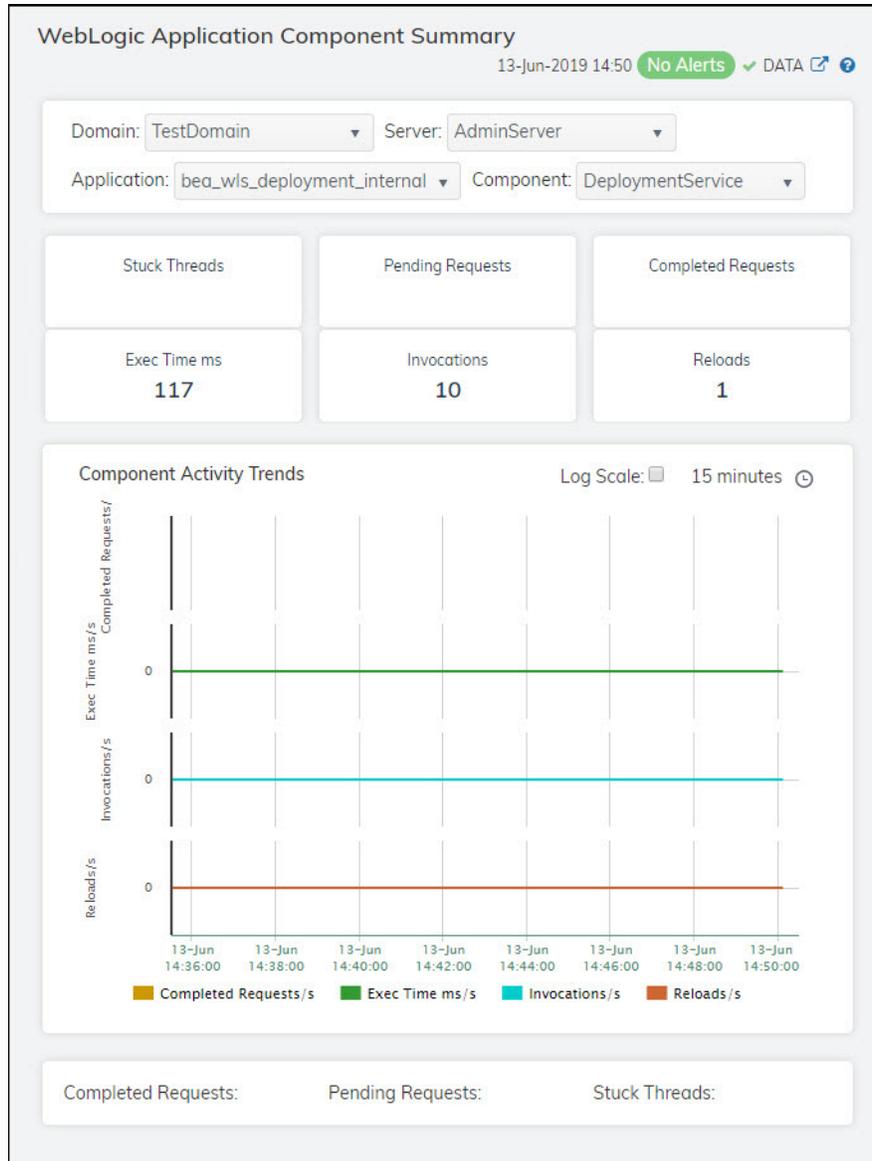
The number of completed requests in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of objects in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

<b>Completed Requests/s</b>	The number of completed requests per second in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of objects in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
<b>Invocations /s</b>	The number of invocations per second in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of objects in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
<b>Invocations</b>	The total number of invocations in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of objects in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
<b>Execution Time ms</b>	The total execution time in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of objects in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
<b>Reloads</b>	The total number of reloads in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of objects in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

## WebLogic Application Component Summary - HTML

Clicking **App Component Summary** in the left/navigation menu opens the **WebLogic Application Component Summary** display, which allows you to view performance, utilization, and trend data for a particular application component on a single WebLogic Server. Hovering over the information boxes at the top of the display provides additional details.

The **Component Activity Trends** trend graph at the bottom of the display traces the completed requests rate, the execution time rate, the invocation creation rate, and the number of reloads per second over a defined time range.







## CHAPTER 7 RTView DataServer for Solace

The RTView DataServer for Solace provides a way to create connections and modify default configuration settings for solution packages and sends collected data to RTViewCentral. RTViewCentral contains the displays associated with the RTView DataServer for Solace which you use to monitor your Solace components. This section describes the HTML version which features an overview display, "[Brokers Overview](#)", and the following Views which can be found under **Components** tab > **Middleware** > **Solace**:

The RTView *DataCollector* for Solace is also available for use with the RTView DataServer for Solace. RTView DataCollector for Solace is used for collecting solution package data and sending it to one or more RTView DataServers. The RTView DataCollector for Solace is useful if you need to distribute data collection.

For an overview and details about configuring RTView Enterprise, including RTViewCentral, RTView DataServers, RTView DataCollectors and solution packages, see the *RTView Enterprise Configuration Guide*.

RTViewCentral contains the following Views and displays that will be populated with data collected via the RTView DataServer for Solace:

- "[Brokers](#)": The displays in this View present message router-level metrics, which show configuration settings, total throughput, current status, errors, and value-added calculations that summarize metrics across all of the VPNs.
- "[CSPF Neighbors](#)": The displays in this View present a topology and metrics of your message routers, VMRs and servers as well as and their configuration settings.
- "[VPNs](#)": The displays in this View present VPN-level metrics.
- "[Clients](#)": The displays in this View present metrics for all clients of the message router. These views can be filtered to limit the displays to clients for a single VPN.
- "[Bridges](#)": The displays in this View present a topology and metrics of your bridges and VPNs. These views can be filtered to limit the displays to bridges for a single VPN.
- "[Endpoints](#)": The displays in this View present metrics for topics and queues on the message router, which can be filtered to limit the displays to topics and queues for a single VPN.
- "[Capacity](#)": The displays in this View present current metrics, alert count and severity at the message router level.
- "[Syslog Events](#)": View all Syslog events for your Solace message routers.

**Note:** This document assumes familiarity with the products monitored. For additional details, refer to vendor documentation.

## Brokers

These displays provide detailed metrics for brokers and their connected brokers. Displays in this View are:

- **"Brokers Overview"**: Health snapshot of top 10 most utilized VPNs, trend graphs trace key performance metrics such as messages sent/received and connected clients.
- **"Brokers Heatmap"**: A color-coded heatmap view of the current status of each of your brokers.
- **"Brokers Table"**: A tabular view of all available broker performance data.
- **"Broker Summary"**: Current and historical metrics for a single broker.
- **"Broker Sensors"**: Provides value and status information for all sensors on a single broker or for all sensors for all brokers.
- **"Broker Provisioning"**: Provides broker details such as host, chassis, redundancy, memory, and fabric data for a particular broker.
- **"Broker Interface"**: Provides detailed data and status information for the interfaces associated with one or all broker(s). You can also view current and historical amounts of incoming and outgoing packets and bytes for a selected interface in a trend graph.
- **"Brokers Message Spool"**: Provides status and usage data for message spools associated with one or all broker(s).

### Brokers Overview

The **Brokers Overview** is the top-level display, which provides a good starting point for immediately getting the status of all your brokers on your Data Server.

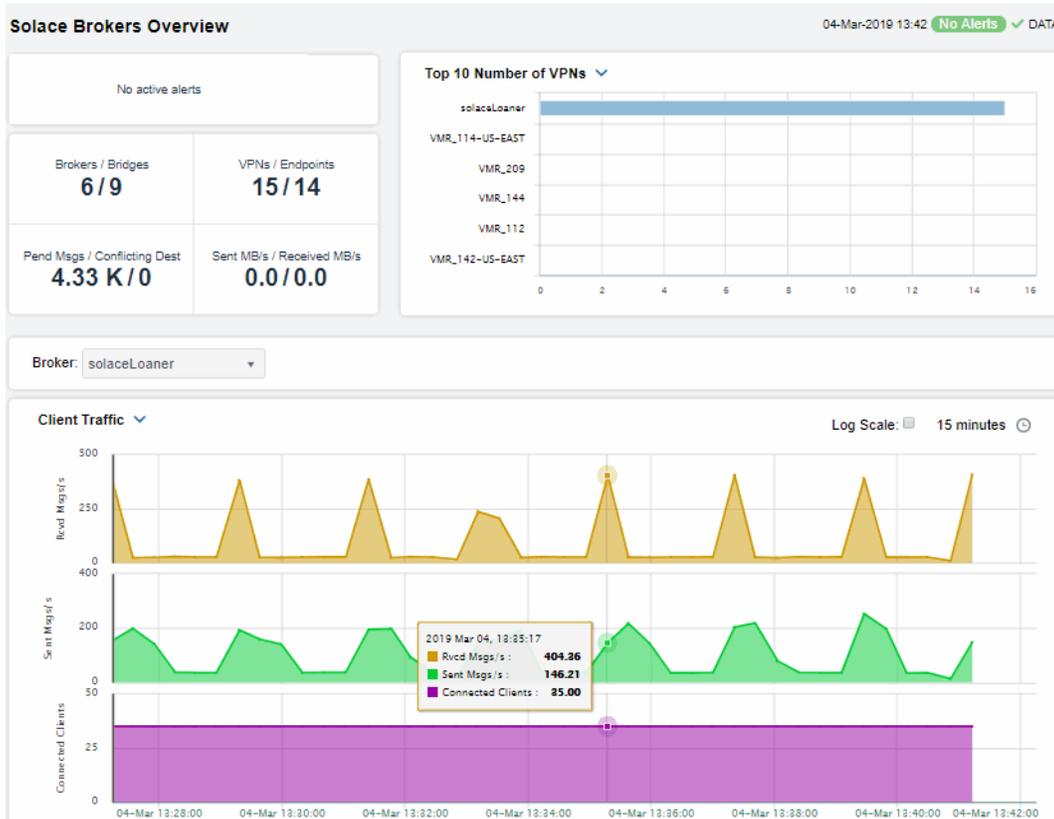
Select a data server, broker and metric from the drop-down menus. Consider keeping this display open for monitoring at a glance. You can easily view the current data for that Data Server including:

- **Top 10** most utilized **VPNs / Endpoints, Clients Connected** and **Spooled Messages**.
- The number of **Brokers / Bridges**.
- The number of **Pending Messages / Conflicting Destinations**.
- The number of **Sent MBs per second / Received MB per second**.

You can hover over each area in the upper half of the Overview to see more detail. You can also drill down to see even more detail by clicking on each metric card in the Overview.

The bottom half of the display provides a performance trend graph for queries for a selected broker. The trend graph traces the performance metric you select: **Client Traffic, Spool Msgs** or **Memory**.

You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.



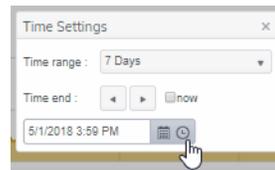
- CRITICAL** Total number of current critical alerts for brokers on the selected data server.
- WARNING** Total number of current warning alerts for brokers on the selected data server.
- Brokers/Bridges** Total number of brokers/bridges on the selected data server.
- VPNs/Endpoints** Total number of VPNs/endpoints on the selected data server.
- Pending Msgs/Conflicting Dest** Total number of pending messages/conflicting destinations on the selected data server.
- Sent MBs/Received MBs** Total number of MBs sent/MBs received on the selected data server.
- Top 10 Number of VPNs** Ten brokers with the greatest number of connected VPNs.
- Broker** Select a broker to trace performance metrics in the trend graph, then choose a metric:
  - Client Traffic:** Traces the number of messages received per second, messages sent per second and the number of connected clients.
  - Spool Msgs:** Traces the number of spooled messages and spool size (in megabytes.)

## Time Settings



By default, the time range end point is the current time. To change the time range, click the **Time Settings** and either:

- choose a **Time range** from 5 Minutes to 7 Days in the drop-down menu.
- specify begin/end dates using the calendar .
- specify begin/end time using the clock .



Toggle forward/backward in the trend graph per the period you choose (from the **Time range** drop-down menu) using arrows



Restore settings to current time by selecting **now** .

## Log Scale

Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.

## Brokers Heatmap

View the current status and alerts in a heatmap of all brokers or a subset of brokers. Use the **Show** dropdown menu to choose **All** brokers, **Expired** brokers, **Unexpired** brokers or only brokers in **Standby** mode

Each rectangle in the heatmap is a single broker where the rectangle size represents the number of connections. The rectangle color maps where the current value is on its color gradient bar. Select a broker from the drop-down menu. For example, by default, **Alert Severity** is shown:

### Alert Severity

The current alert severity. Values range from **0** - **2**, as indicated in the color gradient bar, where **2** is the highest Alert Severity:

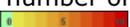
- Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
- Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
- Green indicates that no metrics have exceeded their alert thresholds.

Click a rectangle to drill down to details about a broker in the "**Broker Summary**" display.

Mouse over a rectangle to see additional details. Use the check-box  to include / exclude **Connected** brokers and enable **Log Scale** mode.

Consider keeping this display open for monitoring your Solace brokers at a glance.



<b>Alert Severity</b>	<p>The current alert severity. Values range from <b>0</b> - <b>2</b>, as indicated in the color gradient  bar, where <b>2</b> is the highest Alert Severity:</p> <ul style="list-style-type: none"> <li> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Alert Count</b>	<p>The total number of critical and warning alerts. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.</p>
<b># Msgs Spooled</b>	<p>The total number of spooled messages. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the defined alert threshold of <b>SolMsgRouterPendingMsgsHigh</b>. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>Total Msgs Rcvd</b>	<p>The total number of received messages. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of total messages received in the heatmap. The middle value in the gradient bar indicates the average count.</p>
<b>Total Msgs Sent</b>	<p>The total number of sent messages. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of total messages sent in the heatmap. The middle value in the gradient bar indicates the average count.</p>
<b>Total Msgs/ sec Rcvd</b>	<p>The number of messages received per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the defined alert threshold of <b>SolMsgRouterInboundMsgRateHigh</b>. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>Total Msgs/ sec Sent</b>	<p>The total number of messages sent per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the defined alert threshold of <b>SolMsgRouterOutboundMsgRateHigh</b>. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>Total Bytes/ sec Rcvd</b>	<p>The total number of bytes received per second in the broker. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of <b>SolMsgRouterInboundByteRateHigh</b>. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>Total Bytes/ sec Sent</b>	<p>The total number of bytes sent per second in the broker. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of <b>SolMsgRouterOutboundByteRateHigh</b>. The middle value in the gradient bar indicates the middle value of the range.</p>

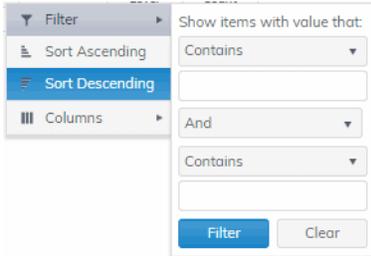
## Brokers Table

Investigate detailed utilization metrics for all brokers. This display provides a tabular view of the performance metrics shown in the "[Brokers Heatmap](#)" (alert level, alert count, and so forth), but with additional metrics such as **Egress** and **Ingress** values.

Use the **Show:** dropdown menu to view the current status of **All** brokers, **Expired** brokers, **Unexpired** brokers or just brokers in **Standby** mode.

Each row in the table contains data for a particular broker. Click a column header to sort column data in ascending or descending order. Double-click on a table row to drill down to the “Broker Summary” display and view metrics for that particular broker. Toggle between the commonly accessed Table and Heatmap displays by clicking the drop down list on the display title.

Search by clicking the right side of a column heading/**Filter** to open the Search, Sort and Choose Columns dialog:



**Brokers:** (in the upper right portion) is the number of brokers in the display.

Use the check-boxes  to include / exclude **Connected** and **Expired** brokers.

**Export to Excel** by right-clicking a column heading.

Toggle between **More Columns** / **Fewer Columns**



Solace Brokers Table 11-Jul-2019 11:57 1 Alert DATA [🔗](#) [📄](#)

Show Connected Only:  Show Expired: All Show Standby: All Brokers: 23

Fewer Columns

Broker	Connected	Alert Level	Alert Count	Host Name	Host Address	VPNs	Total Clients	Total Clieni Connecti
Sol_Mule_Azure			0	solhcdemo0.francecen		0	0	
Sol_Mule_GCP			0	35.234.63.231	35.234.63.231	0	0	
Sol_Mule_SC			0	vmr-mr8v6yiwawhm	34.227.76.129	1	0	
solDemo			0	solace	192.168.220.5	15	40	
Team-2-Ali-cloud			0	47.74.235.254	47.74.235.254	0	0	
Team-2-AWS			0	ec2-35-177-122-45.e	35.177.122.45	0	0	
Team-2-Azure			0	demotion-team20.s		0	0	
Team-2-Google-cloud			0	35.187.64.112	35.187.64.112	0	0	
Team-2-Jab-appliance			0	london.solace.com	217.196.247.77	1	0	
Team-2-Solace-cloud			0	mr-xy4p45t57.messa		0	0	
Team-4-PCF-demo			0	shared-vmr-1.system	35.201.65.176	0	0	
Team-6-Alibaba			0	47.74.235.254	47.74.235.254	0	0	
Team-6-AnalyticsVMR			0	54.179.163.185	54.179.163.185	0	0	
Team-6-OnPremVMR-backup			0	54.191.207.187	54.191.207.187	0	0	
Team-6-OnPremVMR-primar			0	34.214.62.219	34.214.62.219	0	0	
Team-6-PCF			0	shared-vmr-1.system	35.201.65.176	0	0	
Team-6-SolaceCloud			0	mr-91b692durd.messa		0	0	
VMR-112			0	ip-172-30-1-112	172.30.1.112	2	11	

**Column Values**

<b>Broker</b>	The name of the broker.
<b>Connected</b>	<p>The broker state:</p> <ul style="list-style-type: none"> <li><span style="color: red;">●</span> Red indicates that the broker is NOT connected.</li> <li><span style="color: green;">●</span> Green indicates that the broker is connected.</li> </ul>
<b>Alert Severity</b>	<p>The current alert severity:</p> <ul style="list-style-type: none"> <li><span style="color: red;">●</span> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li><span style="color: yellow;">●</span> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li><span style="color: green;">●</span> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Alert Count</b>	The total number of alerts.
<b>Expired</b>	<p>When checked, performance data about the sensor has not been received within the time specified (in seconds) in the <b>\$solRowExpirationTime</b> field in the <b>conf\rtvapm_solmon.properties</b> file. The <b>\$solRowExpirationTimeForDelete</b> field allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response from the sensor. To view/edit the current values, modify the following lines in the <b>.properties</b> file:</p> <pre># Metrics data are considered expired after this number of seconds # collector.sl.rtvview.sub=\$solRowExpirationTime:45 collector.sl.rtvview.sub=\$solRowExpirationTimeForDelete:3600</pre> <p>In the example above, the <b>Expired</b> check box would be checked after 45 seconds, and the row would be removed from the table after 3600 seconds.</p>
<b>Host Name</b>	The name of the host.
<b>Platform</b>	The name of the platform.
<b>OS Version</b>	The version of the operating system.
<b>Up Time</b>	The amount of time that the broker has been up and running.
<b>VPNs</b>	The total number of VPNs configured on the broker.
<b>Total Clients</b>	The total number of clients associated with the broker.
<b>Total Clients Connected</b>	The total number of clients that are currently connected to the broker.
<b>Clients Using Compression</b>	The number of clients who send/receive compressed messages.
<b>Clients Using SSL</b>	The number of clients using SSL for encrypted communications.

<b>Max Client Connections</b>	The maximum number of available client connections.
<b>Endpoints</b>	The total number of Endpoints configured on the broker.
<b>Bridges</b>	The total number of bridges configured on the broker.
<b>Local Bridges</b>	The total number of local bridges configured on the broker.
<b>Remote Bridges</b>	The total number of remote bridges configured on the broker.
<b>Remote Bridge Subscriptions</b>	The total number of remote bridge subscriptions configured on the broker.
<b>Routing Enabled</b>	This check box is checked when the broker is configured to route messages to other brokers.
<b>Routing Interface</b>	The name of the interface configured to support message routing.
<b>Total # Conflicting Destinations</b>	The total number conflicting destinations.
<b>Pending Messages</b>	The number of pending messages on the broker.
<b>Total Client Msgs Rcvd</b>	The total number of client messages received on the broker.
<b>Total Client Msgs Sent</b>	The total number of client messages sent by the broker.
<b>Total Client Msgs Rcvd/sec</b>	The total number of client messages received per second by the broker.
<b>Total Client Msgs Sent/ sec</b>	The total number of client messages sent by the broker.
<b>Total Client Bytes Rcvd</b>	The total number of client bytes received by the broker.
<b>Total Client Bytes Sent</b>	The total number of client bytes sent by the broker.
<b>Total Client Bytes Rcvd/sec</b>	The total number of client bytes received per second by the broker.
<b>Total Client Bytes Sent/sec</b>	The total number of client bytes sent per second by the broker.
<b>Total Client Direct Msgs Rcvd</b>	The total number of direct client messages received by the broker.
<b>Total Client Direct Msgs Sent</b>	The total number of direct client messages sent from the broker.
<b>Total Client Direct Msgs Rcvd/sec</b>	The total number of direct client messages received per second by the broker.
<b>Total Client Direct Msgs Sent/sec</b>	The total number of direct client messages sent per second by the broker.
<b>Total Client Direct Bytes Rcvd</b>	The total number of direct client bytes received by the broker.
<b>Total Client Direct Bytes Sent</b>	The total number of direct client bytes sent by the broker.
<b>Total Client Direct Bytes Rcvd/sec</b>	The total number of direct client bytes received per second by the broker.

<b>Total Client Direct Bytes Sent/sec</b>	The total number of direct client bytes sent per second by the broker.
<b>Total Client Non-Persistent Msgs Rcvd</b>	The total number of non-persistent client messages received by the broker.
<b>Total Client Non-Persistent Msgs Sent</b>	The total number of non-persistent client messages sent by the broker.
<b>Total Client Non-Persistent Msgs Rcvd/sec</b>	The total number of non-persistent client messages received per second by the broker.
<b>Total Client Non-Persistent Msgs Sent/ sec</b>	The total number of non-persistent client messages sent per second by the broker.
<b>Total Client Non-Persistent Bytes Rcvd</b>	The total number of non-persistent client bytes received by the broker.
<b>Total Client Non-Persistent Bytes Sent</b>	The total number of non-persistent client bytes sent by the broker.
<b>Total Client Non-Persistent Bytes Rcvd/sec</b>	The total number of non-persistent client bytes received per second by the broker.
<b>Total Client Non-Persistent Bytes Sent/sec</b>	The total number of non-persistent client bytes sent per second by the broker.
<b>Total Client Persistent Msgs Rcvd</b>	The total number of persistent client messages received by the broker.
<b>Total Client Persistent Msgs Sent</b>	The total number of persistent client messages sent by the broker.
<b>Total Client Persistent Msgs Rcvd/sec</b>	The total number of persistent client messages received per second by the broker.
<b>Total Client Persistent Msgs Sent/ sec</b>	The total number of persistent client messages sent per second by the broker.
<b>Total Client Persistent Bytes Rcvd</b>	The total number of persistent client bytes received by the broker.
<b>Total Client Persistent Bytes Sent</b>	The total number of persistent client bytes sent by the broker.
<b>Total Client Persistent Bytes Rcvd/sec</b>	The total number of persistent client bytes received per second by the broker.
<b>Total Client Persistent Bytes Sent/ sec</b>	The total number of persistent client bytes sent per second by the broker.
<b>Avg Egress Bytes/min</b>	The average number of outgoing bytes per minute.
<b>Avg Egress Compressed Msgs/min</b>	The average number of outgoing compressed messages per minute.
<b>Avg Egress Msgs/min</b>	The average number of outgoing messages per minute.
<b>Avg Egress SSL Msgs/min</b>	The average number of outgoing messages per minute being sent via SSL-encrypted connections.
<b>Avg Egress Uncompressed Msgs/min</b>	The average number of uncompressed outgoing messages per minute.
<b>Avg Ingress Bytes/min</b>	The average number of incoming bytes per minute.
<b>Avg Ingress Compressed Msgs/min</b>	The average number of compressed incoming message per minute.
<b>Avg Ingress Msgs/min</b>	The average number of incoming messages per minute.

<b>Average Ingress SSL Msgs/min</b>	The average number of incoming messages per minute being received via SSL-encrypted connections.
<b>Avg Ingress Uncompressed Msgs/min</b>	The average number of uncompressed messages per minute.
<b>Current Egress Bytes/sec</b>	The current number of outgoing bytes per second.
<b>Current Egress Compressed Msgs/sec</b>	The current number of outgoing compressed messages per second.
<b>Current Egress Msgs/sec</b>	The current number of outgoing messages per second.
<b>Current Egress SSL Msgs/sec</b>	The current number of outgoing messages per second sent via SSL-encrypted connections.
<b>Current Egress Uncompressed Msgs/sec</b>	The current number of outgoing uncompressed messages per second.
<b>Current Ingress Bytes/sec</b>	The current number of incoming bytes per second.
<b>Current Ingress Compressed Msgs/sec</b>	The current number of incoming compressed messages per second.
<b>Current Ingress Msgs/sec</b>	The current number of incoming messages per second.
<b>Current Ingress SSL Msgs/sec</b>	The current number of incoming messages per second received via SSL-encrypted connections.
<b>Current Ingress Uncompressed Msgs/sec</b>	The current number of incoming uncompressed messages per second.
<b>Ingress Comp Ratio</b>	The percentage of incoming messages that are compressed.
<b>Egress Comp Ratio</b>	The percentage of outgoing messages that are compressed.
<b>Egress Compressed Bytes</b>	The number of outgoing compressed bytes.
<b>Egress SSL Bytes</b>	The number of outgoing compressed bytes being sent via SSL-encrypted connections.
<b>Egress Uncompressed Bytes</b>	The number of outgoing uncompressed bytes.
<b>Ingress Compressed Bytes</b>	The number of incoming compressed bytes.
<b>Ingress SSL Bytes</b>	The number of incoming bytes via SSL-encrypted connections.
<b>Ingress Uncompressed Bytes</b>	The number of incoming uncompressed bytes.
<b>Total Egress Discards</b>	The total number of outgoing messages that have been discarded by the broker.
<b>Total Egress Discards/sec</b>	The total number of outgoing messages per second that have been discarded by the broker.
<b>Total Ingress Discards</b>	The total number of incoming messages that have been discarded by the broker.
<b>Total Ingress Discards/sec</b>	The total number of incoming messages per second that have been discarded by the broker.
<b>Client Authorization Failures</b>	The number of failed authorization attempts

<b>Client Connect Failures (ACL)</b>	The number of client connection failures caused because the client was not included in the defined access list.
<b>Subscribe Topic Failures</b>	The number of failed attempts at subscribing to topics.
<b>TCP Fast Retrans Sent</b>	The total number of messages that were retransmitted as a result of TCP Fast Retransmission (one or more messages in a sequence of messages that were not received by their intended party that were sent again).
<b>Memory (KB)</b>	The total available memory (in kilobytes) on the broker.
<b>Memory Free (KB)</b>	The total amount of available memory (in kilobytes) on the broker.
<b>Memory Used (KB)</b>	The total amount of memory used (in kilobytes) on the broker.
<b>Memory Used %</b>	The percentage of total available memory that is currently being used.
<b>Swap (KB)</b>	The total available swap (in kilobytes) on the broker.
<b>Swap Free (KB)</b>	The total amount of available swap (in kilobytes) on the broker.
<b>Swap Used (KB)</b>	The total amount of swap used (in kilobytes) on the broker.
<b>Swap Used %</b>	The percentage of total available swap that is currently being used.
<b>Subscription Mem Total (KB)</b>	The total amount of available memory (in kilobytes) that can be used by queue/topic subscriptions.
<b>Subscription Mem Free (KB)</b>	The current amount of available memory (in kilobytes) that can be used by queue/topic subscriptions.
<b>Subscription Mem Used (KB)</b>	The current amount of memory (in kilobytes) being used by queue/topic subscriptions.
<b>Subscription Mem Used %</b>	The percentage of available memory being used by queue/topic subscriptions.
<b>Chassis Product Number</b>	The product number of the chassis in which the broker is contained.
<b>Chassis Revision</b>	The revision number of the chassis.
<b>Chassis Serial</b>	The serial number of the chassis.
<b>BIOS Version</b>	The basic input/output system used by the chassis.
<b>CPU-1</b>	The name of the central processing unit (CPU 1) used by the broker.
<b>CPU-2</b>	The name of the central processing unit (CPU 2) used by the broker.
<b>Operational Power Supplies</b>	The number of available power supplies that are operational on the chassis.
<b>Power Redundancy Config</b>	The configuration used by the backup broker.
<b>Max # Bridges</b>	The maximum number of bridges allowed on the broker.

<b>Max # Local Bridges</b>	The maximum number of local bridges allowed on the broker.
<b>Max # Remote Bridges</b>	The maximum number of remote bridges allowed on the broker.
<b>Max # Remote Bridge Subscriptions</b>	The maximum number of remote bridge subscriptions allowed on the broker.
<b>Redundancy Config Status</b>	The status of the redundancy configuration.
<b>Redundancy Status</b>	The status of the redundant broker.
<b>Redundancy Mode</b>	Refer to Solace documentation for more information.
<b>Auto-revert</b>	Refer to Solace documentation for more information.
<b>Mate Router Name</b>	If redundancy is configured, this field lists the redundant broker name (mate broker name).
<b>ADB Link Up</b>	This check box is checked if a broker is set up to use guaranteed messaging and an Assured Delivery Blade (ADB) is set up and working correctly.
<b>ADB Hello Up</b>	Refer to Solace documentation for more information.
<b>Pair Primary Status</b>	The primary status of the broker and its redundant (failover) mate.
<b>Pair Backup Status</b>	Refer to Solace documentation for more information.
<b>Expired</b>	<p>When checked, performance data about the broker has not been received within the time specified (in seconds) in the <b>\$solRowExpirationTime</b> field in the <b>conf\rtvapm_solmon.properties</b> file. The <b>\$solRowExpirationTimeForDelete</b> field allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response from the broker. To view/edit the current values, modify the following lines in the <b>.properties</b> file:</p> <pre># Metrics data are considered expired after this number of seconds # collector.sl.rtvview.sub=\$solRowExpirationTime:45 collector.sl.rtvview.sub=\$solRowExpirationTimeForDelete:3600</pre> <p>In the example above, the <b>Expired</b> check box would be checked after 45 seconds, and the row would be removed from the table after 3600 seconds.</p>
<b>Time Stamp</b>	The date and time the row of data was last updated.

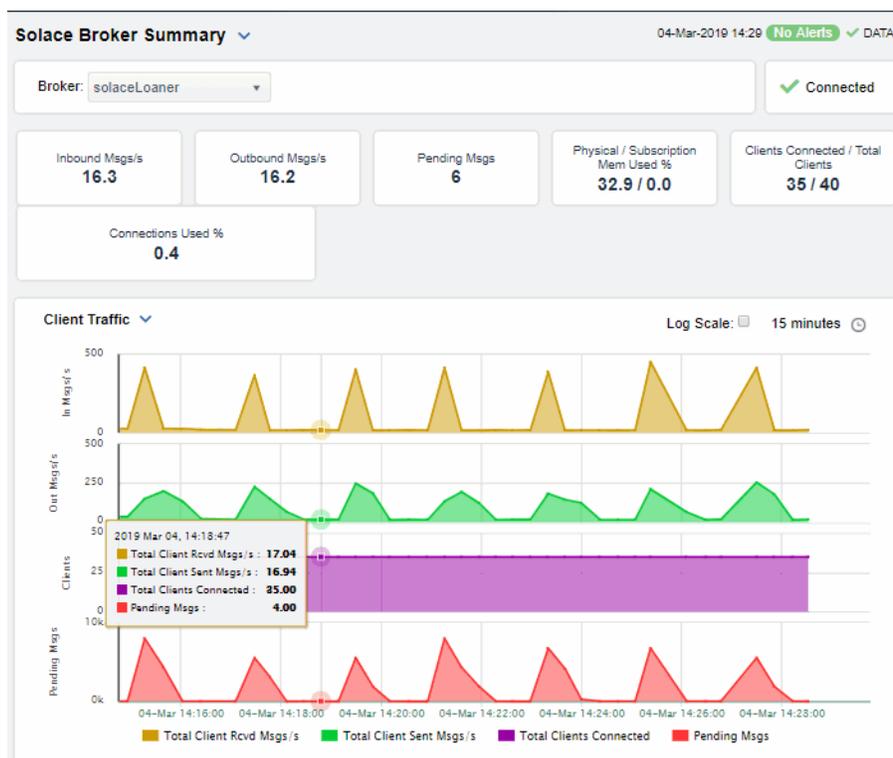
## Broker Summary

View performance and processing details for a single broker, such as the total **Inbound / Outbound Messages per second**, **Pending Messages** and **Clients Connected / Total Clients**.

Choose a broker from the **Broker** drop-down menu to view its total number of connected clients, number of incoming messages, **Up Time**, and additional information. You can also view alert statuses and **Spool Status** data for the broker. You can hover over each area in the upper half of the to see more detail. You can also drill down to see even more detail by clicking on each metric card.

The bottom half of the display provides current and historical performance metrics for the selected broker. The trend graph traces the performance metric you select: **Client Traffic**, **Spool Msgs** or **Memory**.

You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.



The connection status (connected/disconnected).



**Inbound Msgs/s**

The number of messages received per second.

**Outbound Msgs/s**

The number of messages sent per second.

**Pending Msgs/s**

The number of pending messages.

**Physical / Subscription Mem Used %**

The total percentage of physical memory used / the total percentage of subscription memory used.

**Clients Connected / Total Clients**

The current number of clients connected / the total number of clients.

**Connections Used %**

The percentage of connections used.

**Trend Graphs**

Traces the sum for the selected broker.

**Client Traffic**

- **In Msgs/s** - Traces the total number of client messages received per second.
- **Out Msgs/s** - Traces the total number of client messages sent per second.
- **Clients** - Traces the total number of connected clients.
- **Pending Msgs** - Traces the total number of pending messages.

**Spool Msgs**

- **Pending Msgs**- Traces the total number of pending spool messages.
- **Spool Usage MB** - Traces the total amount of space used by spool messages, in megabytes.

**Memory**

- **Memory Used %**- Traces the percent of memory used.

**Subscription Mem Used %** - Traces the percent of memory used by subscriptions.

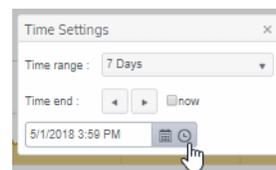
**Log Scale**

Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.

**Time Settings**

By default, the time range end point is the current time. To change the time range, click the **Time Settings**  and either:

- choose a **Time range** from 5 Minutes to 7 Days in the drop-down menu.
- specify begin/end dates using the calendar  .
- specify begin/end time using the clock  .



Toggle forward/backward in the trend graph per the period you choose (from the **Time range** drop-down menu) using arrows   .

Restore settings to current time by selecting **now**

 .

## Broker Sensors

This tabular display contains environmental sensor metrics for a selected broker. Use this display to find out the type, name, value, and status of the sensors.

Select a broker from the drop-down menu. Note that display does not show data for PubSub+ Software as it only applies to brokers.

Search by clicking the right side of a column heading/**Filter** to open the Search, Sort and Choose Columns dialog:

Solace Broker Environmental Sensors							04-Mar-2019 14:38	Alerts	DATA
Broker: solaceLoaner									
Sensor Readings									
Type	Sensor Name	Value	Units	Status	Expired	Time Stamp			
Voltage	BB +1.5V	1.469	volts	OK		04-Mar-2019 14:37:52			
Voltage	BB +1.5V AUX	1.490	volts	OK		04-Mar-2019 14:37:52			
Voltage	BB +1.5V ESB	1.482	volts	OK		04-Mar-2019 14:37:52			
Voltage	BB +1.8V	1.803	volts	OK		04-Mar-2019 14:37:52			
Voltage	BB +12V AUX	12.090	volts	OK		04-Mar-2019 14:37:52			
Voltage	BB +3.3V	3.337	volts	OK		04-Mar-2019 14:37:52			
Voltage	BB +3.3V STB	3.337	volts	OK		04-Mar-2019 14:37:52			
Voltage	BB +5V	5.070	volts	OK		04-Mar-2019 14:37:52			
ThermalMargin	CPU1 Therm Margin	-67.000	degrees C			04-Mar-2019 14:37:52			
ThermalMargin	CPU2 Therm Margin	-59.000	degrees C			04-Mar-2019 14:37:52			
Temperature	Chassis Temp.	23.000	degrees C			04-Mar-2019 14:37:52			
Fan speed	Chassis Fan 1	7714	RPM			04-Mar-2019 14:37:52			
Fan speed	Chassis Fan 2	8057	RPM			04-Mar-2019 14:37:52			
Fan speed	Chassis Fan 3	7714	RPM			04-Mar-2019 14:37:52			
Fan speed	Chassis Fan 4	7543	RPM			04-Mar-2019 14:37:52			
Fan speed	Chassis Fan 5	7371	RPM			04-Mar-2019 14:37:52			
Fan speed	Chassis Fan 6	7288	RPM			04-Mar-2019 14:37:52			
Power system status	Power Redundancy	yes				04-Mar-2019 14:37:52			

### Sensor Readings

Each row in the table is a different sensor on the broker.

<b>Type</b>	See vendor documentation for details.
<b>Sensor Name</b>	The name of the sensor.
<b>Value</b>	Lists the value of the sensor.
<b>Units</b>	Lists the unit of measure for the sensor.
<b>Status</b>	The current status of the sensor.

**Expired**

When checked, performance data about the sensor has not been received within the time specified (in seconds) in the **\$solRowExpirationTime** field in the **conf\rtvapl\_solmon.properties** file. The **\$solRowExpirationTimeForDelete** field allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response from the sensor. To view/edit the current values, modify the following lines in the **.properties** file:

```
# Metrics data are considered expired after this number of seconds
#
collector.sl.rtvapl.sub=$solRowExpirationTime:45
collector.sl.rtvapl.sub=$solRowExpirationTimeForDelete:3600
```

In the example above, the **Expired** check box would be checked after 45 seconds, and the row would be removed from the table after 3600 seconds.

**Time Stamp**

The date and time the row of data was last updated.

**Broker Provisioning**

This display shows provisioning metrics for a single broker. Use this to see the host, platform, chassis, memory, redundancy and fabric data for a specific broker.

Select a broker from the drop-down menus. Search by clicking the right side of a column heading/**Filter** to open the Search, Sort and Choose Columns dialog:

**Broker Provisioning** 04-Mar-2019 14:39 No Alerts DATA

Broker: solaceLoaner

---

Host Name: **solace**                      CPU-1: Intel(R) Xeon(R) CPU E5450 @ 3.00GHz  
 CPU-2: Intel(R) Xeon(R) CPU E5450 @ 3.00GHz    Platform: Solace 3260  
 Chassis Product Number: CHS-3260AC-01-B    BIOS Version: S5000.86B.10.00.0094.101320081858  
 Chassis Revision: 1.4                      Chassis Serial: S009000226                      Power Redundancy Config: 2+2

---

Total Memory (KB): 15,965,652                      Memory Used (KB): 9,332,304                      Memory Used %: 32.89  
 Swap (KB): 2,007,992                      Swap Used (KB): 0                      Swap Used %: 0.0

---

Operational Power Supplies: 4                      Mate Router Name:                      Redundancy Config Status: **Shutdown**  
 Redundancy Status: **Down**                      Redundancy Mode: **N/A**                      Pair Primary Status: **Local Active**  
 Pair Backup Status: **Shutdown**                      Auto Revert: **false**                      ADB Link To Mate Up: **false**  
 ADB Hello To Mate Up: **false**

---

Last Update: 04-Mar-2019 14:39:28

---

**Fabric**

Product	Fw-Version	Card Type	Slot	Serial #
NAB-0801ET-01-A	6.2.0.496	Network Acceleration Blade	1/1	S003000276
		in use by slot 1/1	1/2	
TRB-000000-02-A		Topic Routing Blade	1/3	P004045787
HBA-0204FC-02-A		Host Bus Adapter Blade	1/4	LFC0848B99499
ADB-000000-01-A		Assured Delivery Blade	1/5	S003000844
		empty	2/1	
		empty	2/2	
		empty	2/3	

- Host Name**                      The name of the host.
- Platform**                      The platform on which the broker is running.
- Chassis Product #**                      The product number of the chassis in which the broker is contained.
- Chassis Revision #**                      The revision number of the chassis.
- Chassis Serial #**                      The serial number of the chassis.
- Power Configuration**                      The power configuration used by the chassis.
- Operational Power Supplies**                      The number of available power supplies that are operational on the chassis.
- CPU 1**                      The name of the central processing unit (CPU 1) used by the broker.
- CPU 2**                      The name of the central processing unit (CPU 2) used by the broker.
- BIOS**                      The basic input/output system used by the chassis.
- Memory (KB)**
  - Physical**                      Lists the **Total** amount, the **Free** amount, the **Used** amount, and the **Used %** of physical memory.
  - Swap**                      Lists the **Total** amount, the **Free** amount, the **Used** amount, and the **Used %** of swap memory.
- Redundancy**  
 These fields describe a fault tolerant pair of brokers.

<b>Mate Router Name</b>	If redundancy is configured, this field lists the redundant broker name (mate broker name).
<b>Configuration Status</b>	The status of the configuration for the backup broker.
<b>Redundancy Status</b>	The status of the redundant broker.
<b>Redundancy Mode</b>	Refer to Solace documentation for more information.
<b>Primary Status</b>	The status of the primary broker.
<b>Backup Status</b>	Refer to Solace documentation for more information.
<b>Auto-Revert</b>	Refer to Solace documentation for more information.
<b>ADB Link Up</b>	This check box is checked if a broker is set up to use guaranteed messaging and an Assured Delivery Blade (ADB) is set up and working correctly.
<b>ADB Hello Up</b>	Refer to Solace documentation for more information.

## Fabric

<b>Slot</b>	Displays the slot number on the network switch.
<b>Card Type</b>	The type of card connected to the particular slot.
<b>Product</b>	The product associated with the particular slot.
<b>Serial #</b>	The serial number of the product.
<b>Fw-Version</b>	The firmware version of the product.

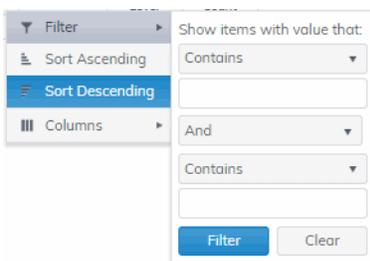
## Broker Interface

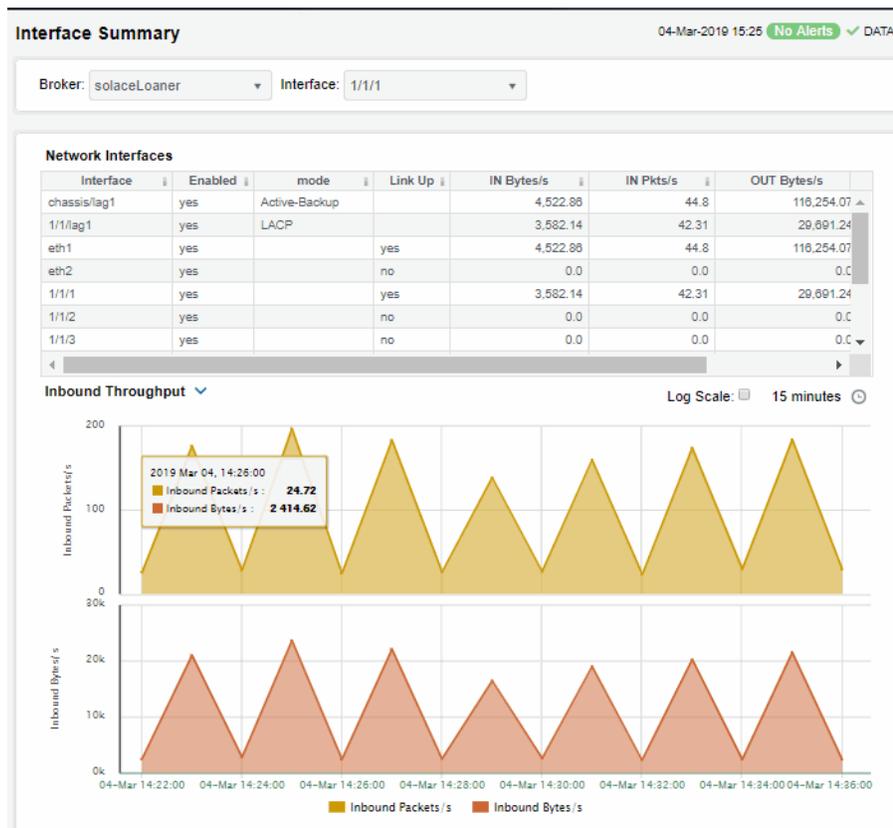
This display lists all network interfaces on a selected broker, and shows network interface status, in/out throughput per second and additional detailed metrics.

Select a broker and interface from the drop-down menus. Each row in the table is a different network interface. Double-click a row to trace its current and historical performance data in the trend graph (bytes in/out and packets in/out per second).

You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.

Search by clicking the right side of a column heading/**Filter** to open the Search, Sort and Choose Columns dialog:





- Interface** The name of the network interface.
- Enabled** Displays whether or not the network interface is enabled.
- mode** Describes how the interface is configured to support networking operations.
- Link Up** Indicates whether the interface is electrically signaling on the transmission medium.
- IN Bytes/sec** The number of bytes per second contained in incoming messages.
- IN Pkts/sec** The number of incoming packets per second.
- OUT Bytes/ sec** The number of bytes per second contained in the outgoing messages.
- OUT Pkts/sec** The number of outgoing packets per second.

**Trend Graphs**

- Inbound Pkts/ sec** Traces the number of incoming packets per second.
- Outbound Bytes/sec** Traces the number of bytes per second contained in the incoming messages.

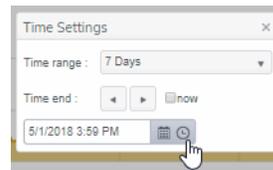
**Log Scale**

Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.

**Time Settings**

By default, the time range end point is the current time. To change the time range, click the **Time Settings**  and either:

- choose a **Time range** from 5 Minutes to 7 Days in the drop-down menu.
- specify begin/end dates using the calendar .
- specify begin/end time using the clock .



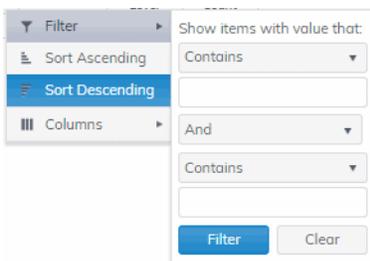
Toggle forward/backward in the trend graph per the period you choose (from the **Time range** drop-down menu) using arrows  .

Restore settings to current time by selecting **now** .

**Brokers Message Spool**

Select a broker from the drop-down menu or select **All**. This display shows operational status and spooling performance metrics (if spooling is enabled on the broker) for one or all brokers.

Search by clicking the right side of a column heading/**Filter** to open the Search, Sort and Choose Columns dialog:



Refer to Solace documentation for details about data in this display.

Pool MB	Msg Spool Used By Queue	Msg Spool Used By DTE	Message Count % Usage	Delivered UnAcked Msgs % Usage	Ingress Flow Count	Ingress Flows Allowed
0.0	13	1	0.0	0.0	18	1

<b>Count</b>	The number of brokers that are using spooling in the table.
<b>Connection</b>	The name of the broker.
<b>Config Status</b>	The status of the connection's configuration.
<b>Operational Status</b>	The operational status of the spool on the broker.
<b>Current Spool Usage (MB)</b>	The current amount of spool used in megabytes on the broker (calculated by summing spool used for each endpoint).
<b>Msg Spool Used By Queue</b>	The amount of spool used by the queue.
<b>Msg Spool Used By DTE</b>	The amount of spool used by DTE.
<b>Message Count % Utilization</b>	The percentage of total messages that use the message spool.
<b>Delivered UnAcked Msgs % Utilization</b>	The percentage of messages delivered via the spool that have not been acknowledged.
<b>Ingress Flow Count</b>	The current incoming flow count.
<b>Ingress Flows Allowed</b>	The total number of incoming flows allowed.
<b>Queue/Topic Subscriptions Used</b>	The number of queue/topic subscriptions used.
<b>Max Queue/Topic Subscriptions</b>	The maximum number of queue/topic subscriptions available.
<b>Sequenced Topics Used</b>	The number of sequenced topics used.
<b>Max Sequenced Topics</b>	The maximum number of sequenced topics available.
<b>Spool Files Used</b>	The number of spool files used.
<b>Spool Files Available</b>	The maximum number of spool files available.

<b>Spool Files % Utilization</b>	The percentage of available spool files that have been used.
<b>Active Disk Partition % Usage</b>	The percentage of available active disk partition that has been used.
<b>Standby Disk Partition % Usage</b>	The percentage of available standby disk partition that has been used.
<b>Disk Usage Current (MB)</b>	The current amount of spool disk usage in megabytes.
<b>Disk Usage Max (MB)</b>	The maximum amount of available spool disk usage in megabytes.
<b>Transacted Sessions Used</b>	The current number of transacted sessions.
<b>Transacted Sessions Max</b>	The maximum number of transacted sessions allowed.
<b>Transacted Session Count % Utilization</b>	The percentage of allowable transacted sessions that have been used.
<b>Transacted Session Resource % Utilization</b>	The percentage of allowable transacted session resources that have been used.
<b>Expired</b>	<p>When checked, performance data about the broker has not been received within the time specified (in seconds) in the <b>\$solRowExpirationTime</b> field in the <b>conf\rtvapm_solmon.properties</b> file. The <b>\$solRowExpirationTimeForDelete</b> field allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response from the broker. To view/edit the current values, modify the following lines in the <b>.properties</b> file:</p> <pre># Metrics data are considered expired after this number of seconds # collector.sl.rtvapm.sub=\$solRowExpiration Time:45 collector.sl.rtvapm.sub=\$solRowExpiration TimeForDelete:3600</pre> <p>In the example above, the <b>Expired</b> check box would be checked after 45 seconds, and the row would be removed from the table after 3600 seconds.</p>

## CSPF Neighbors

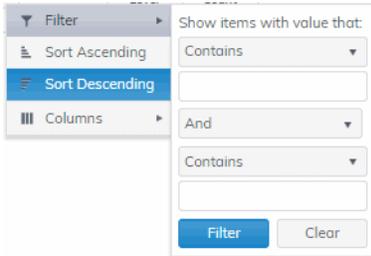
These displays provide detailed data and statuses for CSPF neighbor brokers. You can check trends on network traffic among CSPF neighbors. Note that these displays are empty if you are only monitoring Solace Cloud PubSub+ Brokers. Displays in this View are:

- **"Neighbors Table"**: View metrics for Solace neighbor brokers that use the Content Shortest Path First (CSPF) routing protocol to determine the shortest path in which to send messages from one broker to another broker in the Solace network.
- **"Neighbors Diagram"**: Topological view of CSPF Neighbors that shows broker connections and status of servers (Active/Inactive).
- **"Neighbors Summary"**: View detailed performance metrics for a single Solace neighbor broker that uses the CSPF routing protocol.

## Neighbors Table

This tabular display shows Content Shortest Path First (CSPF) "neighbor" metrics for a broker. Select a broker from the drop-down menu. View metrics for a Solace neighbor broker that uses the CSPF routing protocol to determine the least cost path in which to send messages from one broker to another broker in the Solace network.

Search by clicking the right side of a column heading/**Filter** to open the Search, Sort and Choose Columns dialog:



By default, a subset of available metrics is shown. Use **More Columns/Less Columns** to toggle to the complete set of metrics available (and back to the subset).

Solace CSPF Neighbors Table 14-Aug-2018 16:16 No Alerts DATA [🔗](#) [🔔](#)

Msg Router: - All - More Columns

Show Ok Only:  Show: All Neighbors: 4

Message Router	Name	Expired	State	Sent Msgs/s	Sent Bytes/s	Connections
VMR-112	ip-172-30-1-144		Ok	0.15	24.38	4
VMR-144	ip-172-30-1-112		Ok	0.2	0.0	4
VMR-144	ip-172-30-1-209		Ok	0.2	0.0	4
VMR-209	ip-172-30-1-144		Ok	0.17	29.41	4

- Neighbor Count:** The number of neighbor brokers connected to the selected **Broker**.
- Show:**
  - OK** Select to *only* show neighbor brokers that are connected (**State** is **OK**). By default, this option is not selected (all neighbor brokers are shown).
  - Expired** Select to show *both* expired and non-expired neighbor brokers. By default, this option is not selected (only non-expired neighbor brokers are shown).

**Table:**  
Each table row is a different neighbor broker.

- Broker** The name of the neighbor broker.

<b>State</b>	The current state of the broker.
<b>Up Time</b>	The amount of time the broker has been up and running.
<b>Connections</b>	The number of connections.
<b>Link Cost Actual</b>	Refer to Solace documentation for more information.
<b>Link Cost Configured</b>	Refer to Solace documentation for more information.
<b>Data Port</b>	Refer to Solace documentation for more information.
<b>Expired</b>	<p>When checked, performance data about the broker has not been received within the time specified (in seconds) in the <b>\$solRowExpirationTime</b> field in the <b>conf\rtvadm_solmon.properties</b> file. The <b>\$solRowExpirationTimeForDelete</b> field allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response from the broker. To view/edit the current values, modify the following lines in the <b>.properties</b> file:</p> <pre># Metrics data are considered expired after this number of seconds # collector.sl.rtvadm.sub=\$solRowExpirationTime:45 collector.sl.rtvadm.sub=\$solRowExpirationTimeForDelete:3600</pre> <p>In the example above, the <b>Expired</b> check box would be checked after 45 seconds, and the row would be removed from the table after 3600 seconds.</p>
<b>Timestamp</b>	The date and time the row of data was last updated.

## Neighbors Diagram

Use this topology display to monitor the health of network components: Solace brokers, VMRs and servers. Quickly identify broker neighbors, servers that are inactive and which resources their performance impacts. Drag and drop objects to arrange them on the screen (doing so does not logically impact the Solace brokers, PubSub+ Software and servers).

Each object is a Solace broker, VMR or server. Each are labeled with their name and color coded as follows:

- Red indicates that the object has one or more alerts in a critical state.
- Yellow indicates that the object has one or more alerts in a warning state.
- Green indicates that there are no alerts on the object.
- Gray indicates that the object is off-line.

Mouse-over objects to see their host IP address.

Right-click on VMR objects and select **Open VMR UI** to open the Solace VMR login web page.

**Save:** Saves the arrangement of the objects.

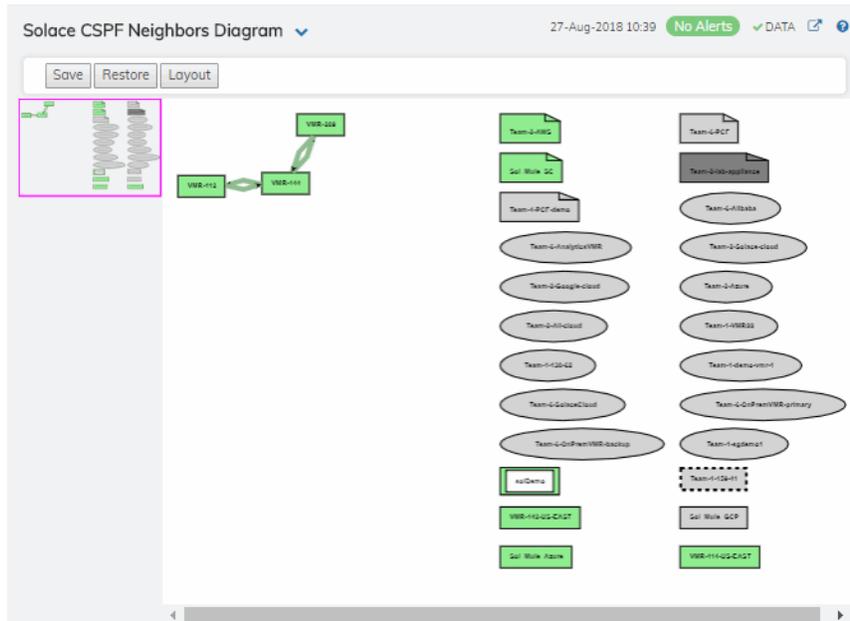
**Restore:** Returns objects to their previous positions.

**Layout:** Toggles between two types of layouts. One layout positions objects to the right so you might scroll in that direction to see them. The other layout pulls the objects close together to the left, vertically and in hierarchical order.

Look at the miniature view in (upper left) to see all objects in either layout. Zoom in on an area in the topology by clicking it in the miniature view.

Drill down to investigate in the ["Neighbors Table"](#).

To monitor network bridges and VPNs, see the ["Bridges - Diagram"](#).



## Neighbors Summary

View neighbor broker current configuration details and message throughput rates.

Select a broker and a neighbor broker from the drop down menus. Check message throughput rates to the neighbor broker, as well as neighbor **Up Time**, **State**, **Data Port**, number of connections and link costs.

You can hover over the metric cards to see more performance metrics and also drill down to see even more detail by clicking on them.

The bottom half of the display provides current and historical performance metrics for the selected broker. The trend graph traces the performance metric you select: **Message Flow** or **Throughput**.

You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.

The trend graph traces the current and historical message throughput (**Data**, **Control**, **Discards** and **Total**).



**Neighbor:** Select the neighbor broker for which you want to show data in the display.

**Connections** The current number of connections.

**Data Msgs/s** Refer to Solace documentation for more information.

**Sent Msgs/s** Refer to Solace documentation for more information.

**Control Msgs/ s** Refer to Solace documentation for more information.

**Data Bytes/s** Refer to Solace documentation for more information.

**Egress Discards/s** The total number of discarded messages sent from the selected **Broker** to the selected **Neighbor** broker since the broker was last started.

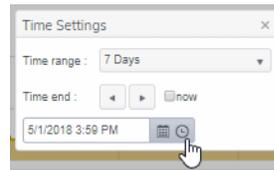
#### Trend Graphs

Traces the rates of messages sent from the selected **Broker** to the selected **Neighbor** broker.

**Sent Msgs/s** Refer to Solace documentation for more information.

**Control Msgs/s** Refer to Solace documentation for more information.

- Discards/s** Traces the number of discarded messages sent, per second, from the selected **Broker** to the selected **Neighbor** broker.
- Log Scale** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Time Settings** By default, the time range end point is the current time. To change the time range, click the **Time Settings**  and either:
- choose a **Time range** from 5 Minutes to 7 Days in the drop-down menu.
  - specify begin/end dates using the calendar .
  - specify begin/end time using the clock .



Toggle forward/backward in the trend graph per the period you choose (from the **Time range** drop-down menu) using arrows   .  
Restore settings to current time by selecting **now** .

## VPNs

You can view data for all VPNs configured on a specific broker in heatmap, table, or grid formats, or you can view data for a single VPN. Displays in this View are:

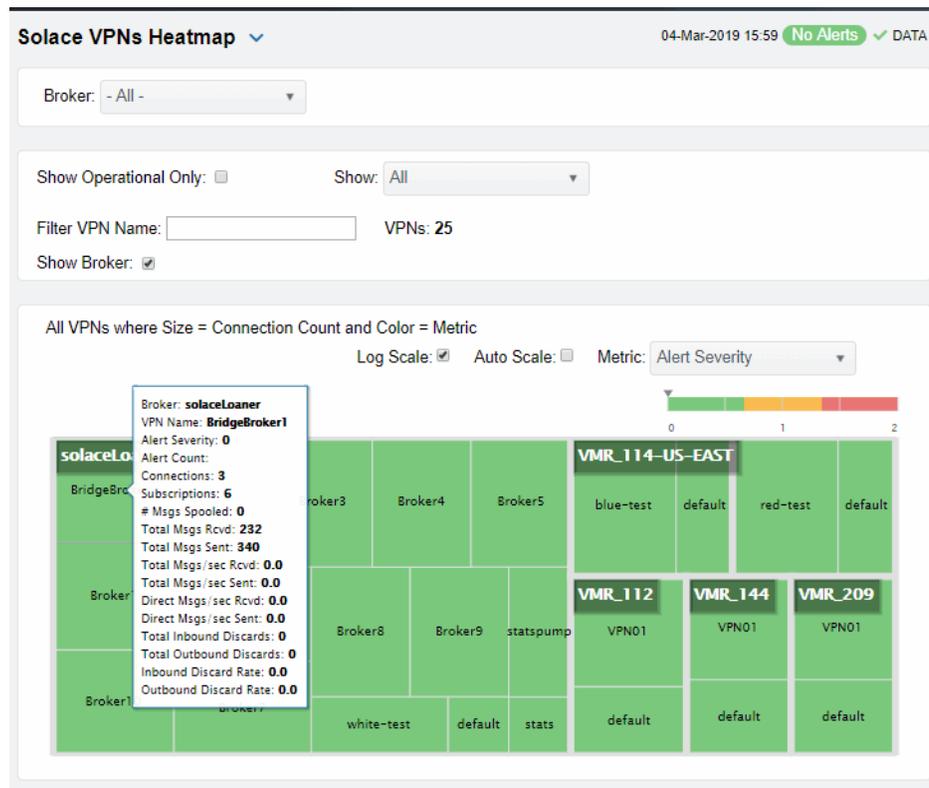
- [“VPNs Heatmap” on page 1008](#): A color-coded heatmap view of the current status of all VPNs configured on a specific broker.
- [“VPNs Table” on page 1012](#): A tabular view of all available data for all VPNs configured on a specific broker.
- [“VPNs Summary” on page 1016](#): Current and historical metrics for a single VPN.

### VPNs Heatmap

View the status of all VPNs configured on a specific broker in a heatmap format, which allows you to quickly identify VPNs with critical alerts. Each rectangle in the heatmap represents a VPN. The rectangle color indicates the alert state and rectangle size represents the number of connections.

Select a broker from the **Broker** drop-down menu, or enter a search string in the **Filter VPN Name** field, and select a metric from the **Metric** drop-down menu. Use the **Show Operational Only** check-box  to include or exclude non-operational VPNs in the heatmap. Use the **Log Scale** and **Auto Scale** check-boxes  to apply log or auto scale. Use the **Show Broker** check-box  to include or exclude broker names in the heatmap.

By default, this display shows **Alert Severity**, but you can mouse over a rectangle to see additional metrics. Drill down and investigate by clicking a rectangle in the heatmap to view details for the selected application in the “VPNs Summary” display.



**Operational** When checked, only shows operational brokers.

**Filter VPN Name** Enter a string to show only VPNs with this string in their name.

**Metric** Choose a metric to view in the display.

#### **Alert Severity**

Visually displays the level at which the VPN has or has not exceeded its alarm level threshold. Values range from **0** - **2**, as indicated in the color gradient bar, where **2** is the highest Alert Severity:

- Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
- Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
- Green indicates that no metrics have exceeded their alert thresholds.

#### **Alert Count**

The total number of critical and warning alerts. The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.

**Connections**

The total number of connections. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **SolVpnConnectionCountHigh**. The middle value in the gradient bar indicates the middle value of the range.

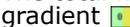
When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

**Subscriptions**

The total number of subscriptions. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **SolVpnSubscriptionCountHigh**. The middle value in the gradient bar indicates the middle value of the range.

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

**# Msgs Spooled**

The total number of spooled messages. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **SolMsgRouterPendingMsgsHigh**. The middle value in the gradient bar indicates the middle value of the range.

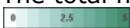
When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

**Total Msgs Rcvd**

The total number of received messages. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of messages received in the heatmap. The middle value in the gradient bar indicates the average count.

The **Auto** flag does not impact this metric.

**Total Msgs Sent**

The total number of sent messages. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of messages sent in the heatmap. The middle value in the gradient bar indicates the average count.

The **Auto** flag does not impact this metric.

**Total Msgs/ sec Rcvd**

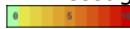
The number of messages received per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **SolVpnInboundMsgRateHigh**. The middle value in the gradient bar indicates the middle value of the range.

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

**Total Msgs/ sec Sent**

The number of messages sent per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **SolVpnOutboundMsgRateHigh**. The middle value in the gradient bar indicates the middle value of the range. When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

**Total Bytes/ sec Rcvd**

The number of bytes contained in messages received per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **SolVpnInboundByteRateHigh**. The middle value in the gradient bar indicates the middle value of the range. When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

**Total Bytes/ sec Sent**

The number of bytes contained in direct messages sent per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **SolMsgRouterOutboundByteRateHigh**. The middle value in the gradient bar indicates the middle value of the range.

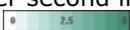
When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

**Direct Msgs/sec Rcvd**

The number of direct messages received per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the average number of direct messages received per second in the heatmap. The middle value in the gradient bar indicates the average count.

The **Auto** flag does not impact this metric.

**Direct Msgs/sec Sent**

The number of direct messages sent per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the average number of direct messages sent per second in the heatmap. The middle value in the gradient bar indicates the average count.

The **Auto** flag does not impact this metric.

**Total Inbound Discards**

The total number of discarded inbound messages in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of discarded inbound messages in the heatmap. The middle value in the gradient bar indicates the average count.

The **Auto** flag does not impact this metric.

<b>Total Outbound Discards</b>	<p>The total number of discarded outbound messages in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of discarded outbound messages in the heatmap. The middle value in the gradient bar indicates the average count.</p> <p>The <b>Auto</b> flag does not impact this metric.</p>
<b>Inbound Discard Rate</b>	<p>The number of discarded inbound messages per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of <b>SolVpnInboundDiscardRateHigh</b>. The middle value in the gradient bar indicates the middle value of the range.</p> <p>When <b>Auto</b> is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>
<b>Outbound Discard Rate</b>	<p>The number of discarded outbound messages per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of <b>SolVpnOutboundDiscardRateHigh</b>. The middle value in the gradient bar indicates the middle value of the range.</p> <p>When <b>Auto</b> is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>

## VPNs Table

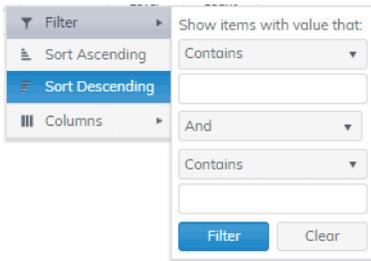
View data shown in the “VPNs Heatmap” display, as well as additional details, in a tabular format. Use this display to view all available data for each VPN associated with a specific broker.

By default, a subset of available metrics is shown. Use **More Columns/Less Columns** to toggle to the complete set of metrics available (and back to the subset).

Select a broker from the **Broker** drop-down menu. Each table row is a different VPN associated with the broker. Click a column header to sort column data in numerical or alphabetical order.

Sort data in numerical or alphabetical order on column headers. Use the check-box  to include / exclude non-operational VPNs. Use the **Show** drop-down to see **All VPNs**, **Expired Only** or **Unexpired Only**. Enter a string in the **Filter VPN Name** field to show only VPNs with this string in their name.

Search by clicking the right side of a column heading/**Filter** to open the Search, Sort and Choose Columns dialog:



Double-click a row to drill down and investigate in the "VPNs Summary" display.

**Solace VPNs Table** 04-Mar-2019 16:05 No Alerts DATA

Broker: - All - Less Columns

Show Operational Only:  Show: All

Filter VPN Name: \* VPNs: 25

Broker	VPN Name	Alert Level	Alert Count	Connections	Operational
solaceLoaner	BridgeBroker1	✓		3	✓
solaceLoaner	Broker1	✓		3	✓
solaceLoaner	Broker10	✓		3	✓
solaceLoaner	Broker2	✓		3	✓
solaceLoaner	Broker3	✓		3	✓
solaceLoaner	Broker4	✓		3	✓
solaceLoaner	Broker5	✓		3	✓
solaceLoaner	Broker6	✓		3	✓
solaceLoaner	Broker7	✓		3	✓
solaceLoaner	Broker8	✓		3	✓
solaceLoaner	Broker9	✓		3	✓
solaceLoaner	default	✓		0	
solaceLoaner	stats	✓		0	
solaceLoaner	statspump	✓		1	✓
solaceLoaner	white-test	✓		1	✓
solaceLoaner	default	✓		1	✓

**Broker**

The name of the broker.

**VPN Name**

The name of the VPN.

**Alert Level**

The maximum level of alerts in the row:

● Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.

● Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.

● Green indicates that no metrics have exceeded their alert thresholds.

**Alert Count**

The total number of active alerts for the VPN.

**Connections**

The total number of connections for the VPN.

<b>Operational</b>	When checked, this status indicates that the VPN is enabled and is operating normally.
<b>Total Unique Subscriptions</b>	The total number of unique subscriptions to the VPN.
<b>Total Client Messages Rcvd</b>	The total number of messages received from clients connected to the VPN.
<b>Total Client Messages Sent</b>	The total number of messages sent to clients connected to the VPN.
<b>Total Client Bytes Rcvd</b>	The total number of bytes contained in messages received from clients connected to the VPN.
<b>Total Client Bytes Sent</b>	The total number of bytes contained in messages sent to clients connected to the VPN.
<b>Total Client Msgs/sec Rcvd</b>	The total number of messages received per second from clients connected to the VPN.
<b>Total Client Msgs /sec Sent</b>	The total number of messages sent per second to clients connected to the VPN.
<b>Total Client Bytes/sec Rcvd</b>	The total number of bytes contained in messages received per second from clients connected to the VPN.
<b>Total Client Bytes/sec Sent</b>	The total number of bytes contained in messages sent per second to clients connected to the VPN.
<b>Client Direct Msgs Rcvd</b>	The total number of direct messages received from clients connected to the VPN.
<b>Client Direct Msgs Sent</b>	The total number of direct messages sent to clients connected to the VPN.
<b>Client Direct Bytes Rcvd</b>	The total number of bytes contained in direct messages received from clients connected to the VPN.
<b>Client Direct Bytes Sent</b>	The total number of bytes contained in direct messages sent to clients connected to the VPN.
<b>Client Direct Msgs/sec Rcvd</b>	The total number of direct messages received per second from clients connected to the VPN.
<b>Client Direct Msgs/sec Sent</b>	The total number of direct messages sent per second to clients connected to the VPN.
<b>Client Direct Bytes/sec Rcvd</b>	The total number of bytes contained in the direct messages received per second from clients connected to the VPN.
<b>Client Direct Bytes/sec Sent</b>	The total number of bytes contained in the direct messages sent per second to clients connected to the VPN.
<b>Client NonPersistent Msgs Rcvd</b>	The total number of non-persistent messages received from clients connected to the VPN.
<b>Client NonPersistent Msgs Sent</b>	The total number of non-persistent messages sent to clients connected to the VPN.
<b>Client NonPersistent Bytes Rcvd</b>	The total number of bytes contained in the non-persistent messages received from clients connected to the VPN.
<b>Client NonPersistent Bytes Sent</b>	The total number of bytes contained in the non-persistent messages sent per second to clients connected to the VPN.
<b>Client NonPersistent Msgs/sec Rcvd</b>	The total number of non-persistent messages received per second from clients connected to the VPN.
<b>Client NonPersistent Msgs/sec Sent</b>	The total number of non-persistent messages sent per second to clients connected to the VPN.

<b>Client NonPersistent Bytes/sec Rcvd</b>	The total number of bytes contained in the non-persistent messages received per second from clients connected to the VPN.
<b>Client NonPersistent Bytes/sec Sent</b>	The total number of bytes contained in the non-persistent messages sent per second to clients connected to the VPN.
<b>Client Persistent Msgs Rcvd</b>	The total number of persistent messages received from clients connected to the VPN.
<b>Client Persistent Msgs Sent</b>	The total number of persistent messages sent to clients connected to the VPN.
<b>Client Persistent Bytes Rcvd</b>	The total number of bytes contained in persistent messages received from clients connected to the VPN.
<b>Client Persistent Bytes Sent</b>	The total number of bytes contained in persistent messages sent to clients connected to the VPN.
<b>Client Persistent Msgs/sec Rcvd</b>	The total number of persistent messages received per second from clients connected to the VPN.
<b>Client Persistent Msgs/sec Sent</b>	The total number of persistent messages sent per second to clients connected to the VPN.
<b>Client Persistent Bytes/sec Rcvd</b>	The total number of bytes contained in the persistent messages received per second from clients connected to the VPN.
<b>Client Persistent Bytes/sec Sent</b>	The total number of bytes contained in the persistent messages sent per second to clients connected to the VPN.
<b>Total In Discards</b>	The total number of discarded incoming messages.
<b>Total In Discards/sec</b>	The number of discarded incoming messages per second.
<b>Total Out Discards</b>	The total number of discarded outgoing messages.
<b>Total Out Discards/sec</b>	The number of discarded outgoing messages per second.
<b>Max Spool Usage (MB)</b>	The maximum amount of disk storage (in megabytes) that can be consumed by all spooled message on the VPN.
<b>Authentication Type</b>	The defined authentication type on the VPN.

**Expired**

When checked, performance data about the VPN has not been received within the time specified (in seconds) in the **\$solRowExpirationTime** field in the **conf\rtvapm\_solmon.properties** file. The **\$solRowExpirationTimeForDelete** field allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response from the VPN. To view/edit the current values, modify the following lines in the **.properties** file:

```
# Metrics data are considered expired after this
number of seconds
#
collector.sl.rtvview.sub=$solRowExpirationTime:4
5
collector.sl.rtvview.sub=$solRowExpirationTimeFo
rDelete:3600
```

In the example above, the **Expired** check box would be checked after 45 seconds, and the row would be removed from the table after 3600 seconds.

**Time Stamp**

The date and time the row data was last updated.

**VPNs Summary**

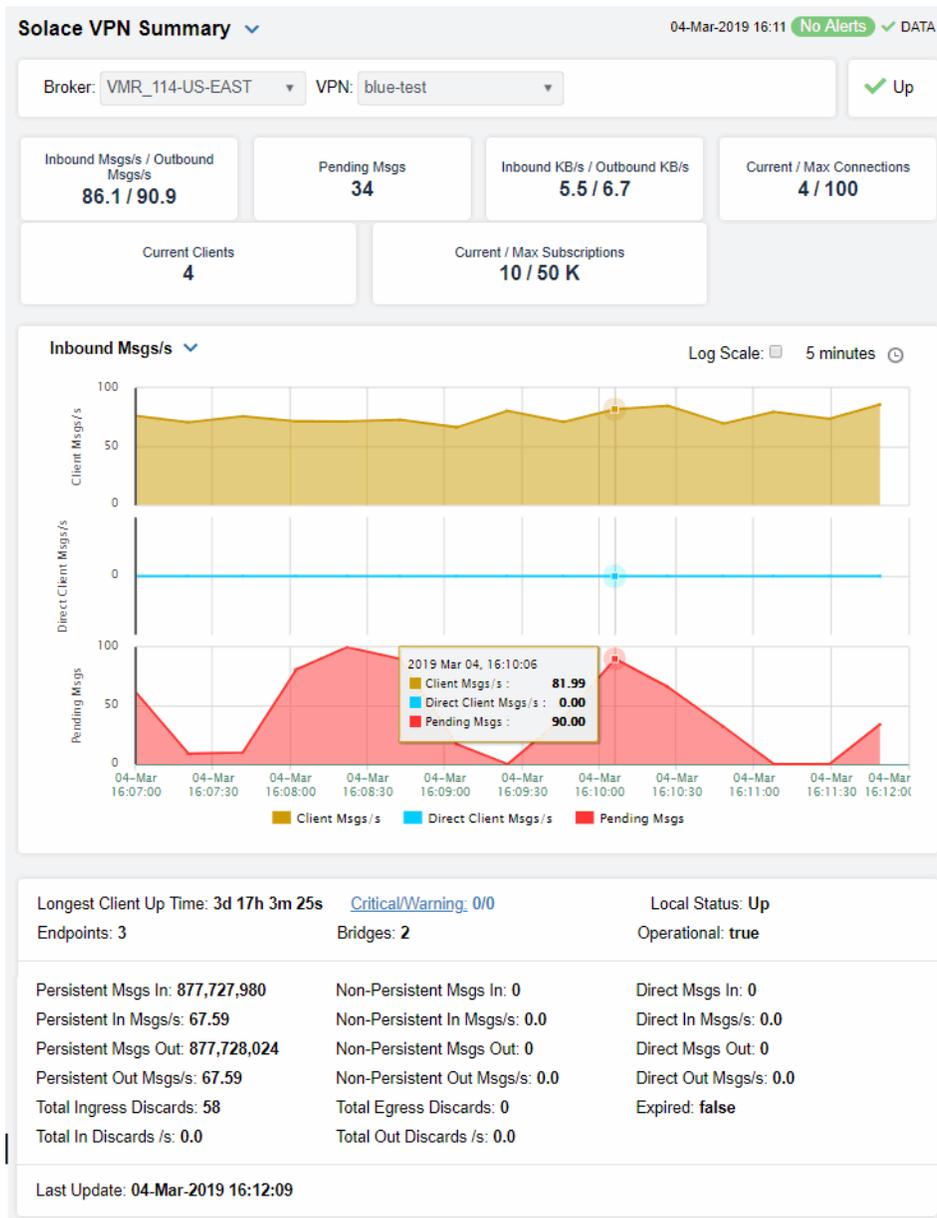
View neighbor broker current configuration details and message throughput rates.

Select a broker and a neighbor broker from the drop down menus. Check message throughput rates to the neighbor broker, as well as neighbor **Up Time, State, Data Port**, number of connections and link costs.

You can hover over the metric cards to see more performance metrics and also drill down to see even more detail by clicking on them.

The bottom half of the display provides current and historical performance metrics for the selected broker. The trend graph traces the performance metric you select: **Ingress Flows, Egress Flows** or **Spool Msgs**.

You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.



**Alerts**

- Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
- Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
- Green indicates that no metrics have exceeded their alert thresholds.

**Up**

**Inbound/Outbound Msgs/s** The number of inbound/outbound messages per second.

- Pending Msgs**                      The number of pending messages.
- Inbound/Outbound KB/s**        The number of inbound/outbound messages in KBs per second.
- Current/Max Connections**      The total number of current connections / maximum number of supported connections for the VPN.
- Current Clients**                 The number of connected clients.
- Current/Max Subscriptions**    The total number of current subscribers and maximum number of supported subscribers for the VPN.

**Inbound Msgs/s Trend Graphs**

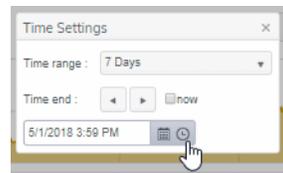
Traces the sum of inbound message processing for the selected VPN.

- **Pending Msgs:** The number of pending messages for the VPN.
- **Client Msgs/sec:** The rate of incoming messages (per second) from client.
- **Direct Client Msgs/sec:** The rate of direct incoming messages (per second) from the direct client.

**Log Scale**                              Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.

**Time Settings**                        By default, the time range end point is the current time. To change the time range, click the **Time Settings**  and either:

- choose a **Time range** from 5 Minutes to 7 Days in the drop-down menu.
- specify begin/end dates using the calendar  .
- specify begin/end time using the clock  .



Toggle forward/backward in the trend graph per the period you choose (from the **Time range** drop-down menu) using arrows   .

Restore settings to current time by selecting **now**  .

- Longest Client Up Time**        The number of days, hours and minutes for the longest, currently active, client connection.
- Endpoints**                         The number of endpoints.
- Persistent Msgs In**                The total number of incoming persistent messages.
- Persistent In Msgs/s**              The number of incoming persistent messages per second.
- Persistent Msgs Out**              The total number of outgoing persistent messages.
- Persistent Out Msgs/s**            The number of outgoing persistent messages per second.
- Total In Discards**                The total number of incoming messages that were discarded.

<b>Total In Discards/sec</b>	The total number of incoming messages that were discarded, per second.
<b>Critical/Warning</b>	The number of critical alerts / warning alerts which also opens the <b>Alerts Table</b> .
<b>Bridges</b>	The number of bridges.
<b>Non-Persistent Msgs In</b>	The total number of incoming non-persistent messages.
<b>Non-Persistent In Msgs/s</b>	The number of incoming non-persistent messages per second.
<b>Non-Persistent Msgs Out</b>	The total number of outgoing non-persistent messages.
<b>Non-Persistent Out Msgs/s</b>	The number of outgoing non-existent messages per second.
<b>Total Out Discards</b>	The total number of outgoing messages that were discarded.
<b>Total Out Discards/sec</b>	The total number of outgoing messages that were discarded, per second.
<b>Direct Msgs In</b>	The total number of incoming direct messages.
<b>Direct In Msgs/s</b>	The number of incoming direct messages per second.
<b>Direct Msgs Out</b>	The total number of outgoing direct messages.
<b>Direct Out Msgs/s</b>	The number of outgoing direct messages per second.
<b>Expired</b>	<p>When <b>true</b>, performance data about the VPN has not been received within the time specified (in seconds) in the <b>\$solRowExpirationTime</b> field in the <b>conf\rtvapm_solmon.properties</b> file. The <b>\$solRowExpirationTimeForDelete</b> field allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response from the VPN. To view/edit the current values, modify the following lines in the <b>.properties</b> file:</p> <pre># Metrics data are considered expired after this number of seconds # collector.sl.rtvview.sub=\$solRowExpirationTime:45 collector.sl.rtvview.sub=\$solRowExpirationTimeForDelete:3600</pre> <p>In the example above, the <b>Expired</b> check box would be checked after 45 seconds, and the row would be removed from the table after 3600 seconds.</p>
<b>Last Update</b>	The date and time of the last data update.

## Clients

These displays allow you to view the current and historical metrics for clients configured on a VPN. Displays in this View are:

- **"Clients Table"**: A tabular view of data for all clients configured on a VPN.
- **"Client Summary"**: Current and historical metrics for a single client configured on a VPN.

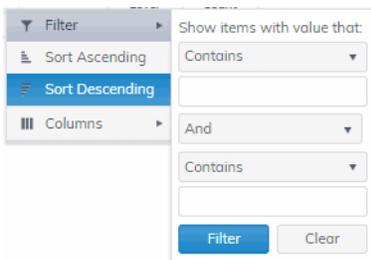
## Clients Table

View VPN clients configured on all brokers, a single broker, all VPNs or a single VPN. Each table row is a different VPN client connection. Use the drop-down menus to show **All**, **Expired** or **Unexpired** clients as well as **All**, **Internal** or **Primary** clients (processes that run on the broker under the Solace OS). Enter a string for **Filter Client Name** to show only clients with this string in their name.

By default, a subset of available metrics is shown. Use **More Columns/Less Columns** to toggle to the complete set of metrics available (and back to the subset).

This display is populated by two caches, SolClientsStats and SolClients. SolClientsStats provides most of the data. SolClients provides the static data. If the SolClients cache encounters an issue the static fields in this display are blank.

Search by clicking the right side of a column heading/**Filter** to open the Search, Sort and Choose Columns dialog:



Double-click a row to drill down and investigate in the “Client Summary” display.

**Solace Broker Clients Table** 05-Mar-2019 08:38 No Alerts DATA

Broker: - All - VPN: - All - Less Columns

Show Type: All Show: All Filter Client Name:

Clients: 101

Broker	VPN	Client Name	Alert Level	Alert Count	Stc Subc
solaceLoaner	BridgeBroker1	#bridge/local/B114toSolDemo/solace/8798/16	✓		
solaceLoaner	BridgeBroker1	#bridge/local/B142toSolDemo/solace/8798/15	✓		
solaceLoaner	BridgeBroker1	#client	✓		
solaceLoaner	Broker1	#bridge/local/testBridgeToNoWhere/solace/8798/14	✓		
solaceLoaner	Broker1	#bridge/remote/B1_to_B2/vr/solace/8796/0	✓		
solaceLoaner	Broker1	#client	✓		
solaceLoaner	Broker1	S-HOST10/5236/#00010001	✓		
solaceLoaner	Broker10	#bridge/local/B112toSolDemo/solace/8798/14	✓		
solaceLoaner	Broker10	#bridge/local/B144toSolDemo/solace/8798/12	✓		
solaceLoaner	Broker10	#bridge/local/B209toSolDemo/solace/8798/13	✓		
solaceLoaner	Broker10	#client	✓		
solaceLoaner	Broker10	S-HOST10/5152/#00010001	✓		
solaceLoaner	Broker10	S-HOST10/5448/#00010001	✓		
solaceLoaner	Broker2	#bridge/local/B1_to_B2/solace/8798/10	✓		
solaceLoaner	Broker2	#client	✓		
solaceLoaner	Broker2	S-HOST10/5212/#00010001	✓		
solaceLoaner	Broker3	#bridge/local/Bridge_loanerToVMR144/solace/8798/11	✓		

Page 1 of 3 1 - 40 of 101 items

<b>Broker</b>	Lists the name of the selected broker.
<b>VPN</b>	Lists the name of the selected VPN.
<b>Client Name</b>	The name of the client.
<b>Alert Level</b>	The maximum level of alerts in the row:  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
<b>Alert Count</b>	Total number of alerts for the client.
<b>Slow Subscriber</b>	This check box will be checked if the client consistently fails to consume their messages at the offered rate (which causes their egress queues to fill up).
<b>Total Egress Flows</b>	The total number of outgoing flows.
<b>Total Ingress Flows</b>	The total number of incoming flows.
<b>Subscriptions</b>	The total number of subscriptions.
<b>Subscription Msgs Rcvd</b>	The total number of messages received from subscriptions.
<b>Subscription Msgs Sent</b>	The total number of messages sent from subscriptions.
<b>Type</b>	Lists the type of alert.
<b>Uptime</b>	Lists the amount of time the client has been up and running.
<b>Client ID</b>	Lists the client ID.
<b>Client UserName</b>	Lists the user name for the client.
<b>Client Address</b>	The IP Address of the client.
<b>Profile</b>	The client profile that is assigned to the client.
<b>ACL Profile</b>	The access control list profile to which the client is assigned.
<b>Description</b>	Lists a description of the client.
<b>Platform</b>	Lists the platform of the client.
<b>Software Version</b>	The version of the platform.
<b>Total Flows Out</b>	The total number of outbound message flows for the client.
<b>Total Flows In</b>	The total number of inbound message flows for the client.
<b># Subscriptions</b>	The number of subscribers connected to the client.
<b>Add Sub Msgs Rcvd</b>	The number of Add Subscription messages received.
<b>Add Sub Msgs Sent</b>	The number of Add Subscription Messages sent.

<b>Already Exists Msgs Sent</b>	Refer to Solace documentation for more information.
<b>Assured Ctrl Msgs Rcvd</b>	Refer to Solace documentation for more information.
<b>Assured Ctrl Msgs Sent</b>	Refer to Solace documentation for more information.
<b>Total Client Msgs Rcvd</b>	The total number of messages received by the client.
<b>Total Client Msgs Sent</b>	The total number of messages sent by the client.
<b>Total Client Bytes Rcvd</b>	The total number of bytes contained within the messages received by the client.
<b>Total Client Bytes Sent</b>	The total number of bytes contained within the messages sent by the client.
<b>Total Client Msgs Rcvd/sec</b>	The total number of messages received per second by the client.
<b>Total Client Msgs Sent/sec</b>	The total number of messages sent per second by the client.
<b>Total Client Bytes Rcvd/sec</b>	The total number of bytes contained within the messages received per second by the client.
<b>Total Client Bytes Sent/sec</b>	The total number of bytes contained within the messages sent per second by the client.
<b>Ctl Bytes Rcvd</b>	The number of control data bytes received by the client.
<b>CTL Bytes Sent</b>	The number of control data bytes sent by the client.
<b>Ctl Msgs Rcvd</b>	The number of control data messages received by the client.
<b>Ctl Msgs Sent</b>	The number of control data messages sent by the client.
<b>Client Data Bytes Rcvd</b>	The number of bytes contained within the data messages received by the client.
<b>Client Data Bytes Sent</b>	The number of bytes contained within the data messages sent by the client.
<b>Client Data Msgs Rcvd</b>	The number of data messages received by the client.
<b>Client Data Msgs Sent</b>	The number of data messages sent by the client.
<b>Client Direct Msgs Rcvd</b>	The number of direct messages received by the client.
<b>Client Direct Msgs Sent</b>	The number of direct messages sent by the client.
<b>Client Direct Bytes Rcvd</b>	The number of bytes contained within direct messages received by the client.
<b>Client Direct Bytes Sent</b>	The number of bytes contained within direct messages sent by the client.
<b>Client Direct Msgs Rcvd/sec</b>	The number of direct messages received per second by the client.
<b>Client Direct Msgs Sent/sec</b>	The number of direct messages sent per second by the client.
<b>Client Direct Bytes Rcvd/sec</b>	The number of bytes contained within the messages received per second by the client.
<b>Client Direct Bytes Sent/sec</b>	The number of bytes contained within the messages sent per second by the client.
<b>Client NonPersistent Msgs Rcvd</b>	The number of non-persistent messages received by the client.

<b>Client NonPersistent Msgs Sent</b>	The number of non-persistent messages sent by the client.
<b>Client NonPersistent Bytes Rcvd</b>	The number of bytes contained within the non-persistent messages received by the client.
<b>Client NonPersistent Bytes Sent</b>	The number of bytes contained within the non-persistent messages sent by the client.
<b>Client NonPersistent Msgs Rcvd/sec</b>	The number of non-persistent messages received per second by the client.
<b>Client NonPersistent Msgs Sent/sec</b>	The number of non-persistent messages sent per second by the client.
<b>Client NonPersistent Bytes Rcvd/sec</b>	The number of bytes contained within the non-persistent messages received per second by the client
<b>Client NonPersistent Bytes Sent/sec</b>	The number of bytes contained within the non-persistent messages sent per second by the client
<b>Client Persistent Msgs Rcvd</b>	The number of persistent messages received by the client.
<b>Client Persistent Msgs Sent</b>	The number of persistent messages sent by the client.
<b>Client Persistent Bytes Rcvd</b>	The number of bytes contained within the persistent messages received by the client.
<b>Client Persistent Bytes Sent</b>	The number of bytes contained within the persistent messages sent by the client.
<b>Client Persistent Msgs Rcvd/sec</b>	The number of persistent messages received per second by the client.
<b>Client Persistent Msgs Sent/sec</b>	The number of persistent messages sent per second by the client.
<b>Client Persistent Bytes Rcvd/sec</b>	The number of bytes contained within the persistent messages received per second by the client.
<b>Client Persistent Bytes Sent/sec</b>	The number of bytes contained within the persistent messages sent per second by the client.
<b>Denied Dup Clients</b>	Refer to Solace documentation for more information.
<b>Denied Subscribe Permission</b>	The number of denied subscription requests due to improper permissions.
<b>Denied Subscribe Topic-ACL</b>	The number of denied subscriptions to topics due to the fact that the client requesting was not on the Access Control List.
<b>Denied Unsubscribe Permission</b>	The number of denied unsubscribe requests due to improper permissions.
<b>Denied Unsubscribe Topic-ACL</b>	The number of denied unsubscribe requests to topics due to the fact that the client requesting was not on the Access Control List.
<b>DTO Msgs Rcvd</b>	The number of Deliver-To-One messages received by the client.
<b>Egress Compressed Bytes</b>	The number of compressed bytes contained within outgoing messages.
<b>Ingress Compressed Bytes</b>	The number of compressed bytes contained within incoming messages.
<b>Total Ingress Discards</b>	The total number of discarded incoming messages.
<b>Total Egress Discards</b>	The total number of discarded outgoing messages.
<b>Total Ingress Discards/sec</b>	The total number of discarded incoming messages per second.

<b>Total Egress Discards/sec</b>	The total number of discarded outgoing messages per second.
<b>Keepalive Msgs Rcvd</b>	The number of Keepalive messages received by the client.
<b>Keepalive Msgs Sent</b>	The number of Keepalive messages sent by the client.
<b>Large Msgs Rcvd</b>	The number of large messages received by the client.
<b>Login Msgs Rcvd</b>	The number of login message received by the client.
<b>Max Exceeded Msgs Sent</b>	The number of responses sent by the client informing the connected broker(s) that the number of the message(s) sent exceeded the maximum allowed.
<b>Not Enough Space Msgs Sent</b>	The number of responses sent by the client informing the connected broker(s) that the size of the message(s) sent exceeded the maximum allowable size, or that the message caused the client's Local Spool Quota to exceed the maximum amount of space.
<b>Not Found Msgs Sent</b>	Refer to Solace documentation for more information.
<b>Parse Error on Add Msgs Sent</b>	Refer to Solace documentation for more information.
<b>Parse Error on Remove Msgs Sent</b>	Refer to Solace documentation for more information.
<b>Remove Subscription Msgs Rcvd</b>	The number of remove subscription requests received by the client.
<b>Remove Subscription Msgs Sent</b>	The number of remove subscription requests sent by the client.
<b>Subscribe Client Not Found</b>	The number of subscription requests for clients that were not found.
<b>Unsubscribe Client Not Found</b>	The number of unsubscribe requests for clients that were not found.
<b>Update Msgs Rcvd</b>	Refer to Solace documentation for more information.
<b>Update Msgs Sent</b>	Refer to Solace documentation for more information.
<b>Expired</b>	<p>When checked, performance data about the client has not been received within the time specified (in seconds) in the <b>\$solRowExpirationTime</b> field in the <b>conf\rtvapl_solmon.properties</b> file. The <b>\$solRowExpirationTimeForDelete</b> field allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response from the client. To view/edit the current values, modify the following lines in the <b>.properties</b> file:</p> <pre># Metrics data are considered expired after this number of seconds # collector.sl.rtvapl.sub=\$solRowExpirationTime: 45 collector.sl.rtvapl.sub=\$solRowExpirationTimeF orDelete:3600</pre> <p>In the example above, the <b>Expired</b> check box would be checked after 45 seconds, and the row would be removed from the table after 3600 seconds.</p>
<b>Timestamp</b>	The date and time the row of data was last updated.

## Client Summary

View current and historical performance and utilization metrics for a single VPN client.

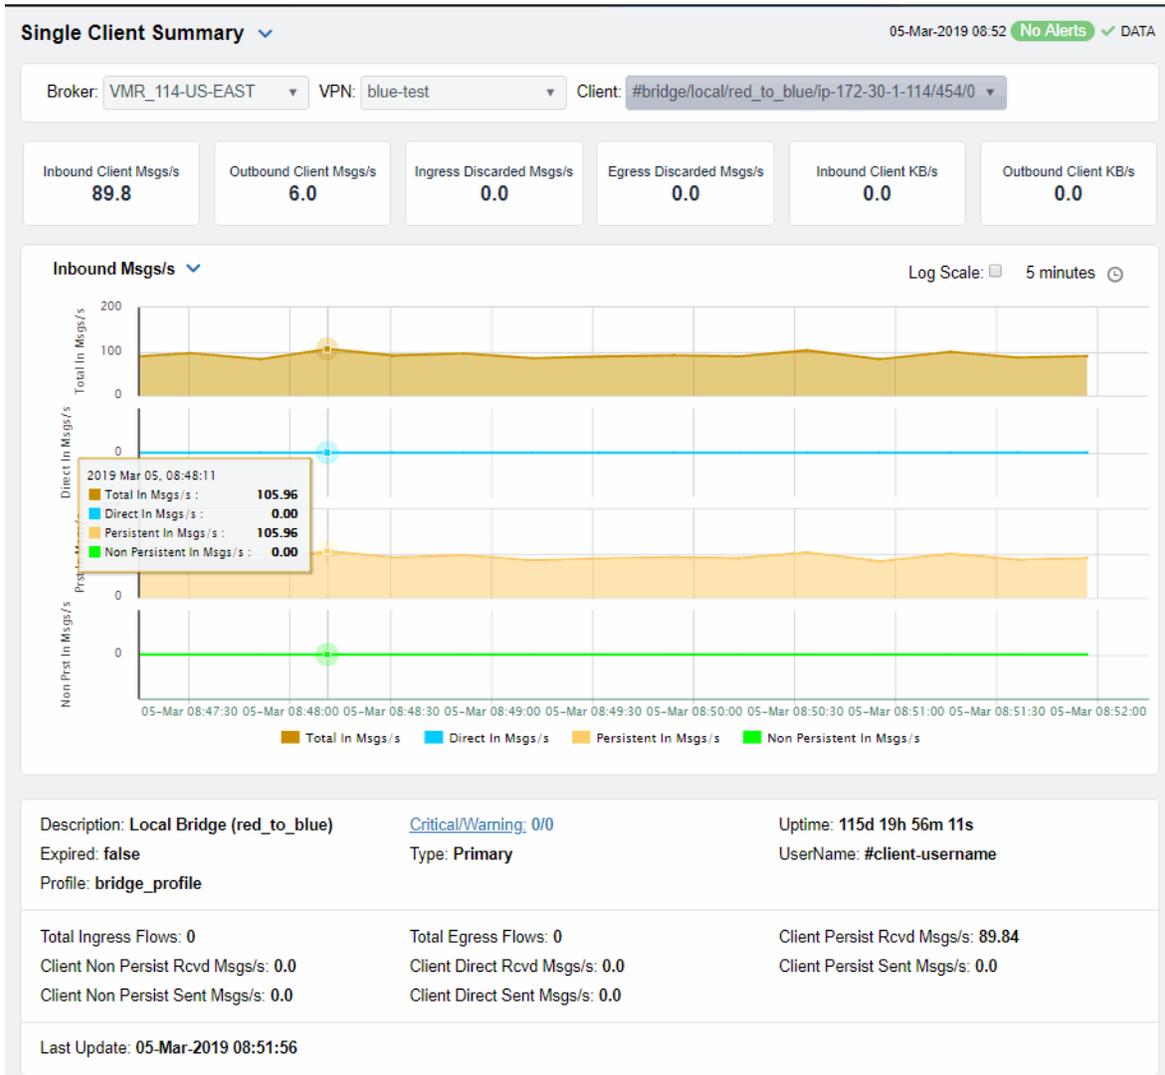
Select a broker, VPN and client from the drop-down menus. You can view the **Client Type**, the **User Name**, the **Client ID**, the associated **Platform**, the current **Up Time**, and additional information specific to the client. You can also view the total number of incoming and outgoing messages, as well as the number of incoming and outgoing persistent, non-persistent, direct, and discarded messages.

You can hover over the metric cards to see more performance metrics and also drill down to see even more detail by clicking on them.

The bottom half of the display provides current and historical performance metrics for the selected broker. The trend graph traces the performance metric you select: **Ingress Flows** or **Egress Flows**.

You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.

This display is populated by two caches, SolClientsStats and SolClients. SolClientsStats provides most of the data. SolClients provides the static data. If the SolClients cache encounters an issue the graphic elements that have no data are replaced with **N/A**.



- Inbound Client Msgs /sec**      The number of incoming client messages per second.
- Outbound Client Msgs /sec**      The number of outgoing client messages per second.
- Ingress Discarded Msgs /sec**      The number of discarded ingress messages per second.
- Egress Discarded Msgs /sec**      The number of discarded egress messages per second.
- Inbound Client KB/sec**      The amount of incoming data from the client in KBs per second.
- Outbound Client KB/sec**      The amount of outgoing data for the client in KBs per second.

**Trend Graphs**

Traces the sum of message processing for the selected client.

- **Total In Msgs/sec**: The number of incoming messages (per second) for the client.
- **Dir-In Msgs/sec**: The number of incoming direct messages (per second) for the client.
- **Persistent In Msgs/sec**: The number of incoming persistent messages (per second) for the client.
- **Non Persistent In Msgs/sec**: The number of incoming non-persistent messages (per second) for the client.

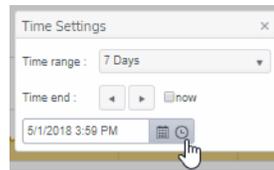
**Log Scale**

Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.

**Time Settings**

By default, the time range end point is the current time. To change the time range, click the **Time Settings**  and either:

- choose a **Time range** from 5 Minutes to 7 Days in the drop-down menu.
- specify begin/end dates using the calendar  .
- specify begin/end time using the clock  .



Toggle forward/backward in the trend graph per the period you choose (from the **Time range** drop-down menu) using arrows



Restore settings to current time by selecting **now**  .

**Description**

The description of the client.

**Expired**

When checked, performance data about the client has not been received within the time specified (in seconds) in the **\$solRowExpirationTime** field in the **conf\rtvapm\_solmon.properties** file. The **\$solRowExpirationTimeForDelete** field allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response from the client. To view/edit the current values, modify the following lines in the **.properties** file:

```
# Metrics data are considered expired after this number of
seconds
#
collector.sl.rtview.sub=$solRowExpirationTime:45
collector.sl.rtview.sub=$solRowExpirationTimeForDelete:3600
```

In the example above, the **Expired** check box would be checked after 45 seconds, and the row would be removed from the table after 3600 seconds.

**Profile**

The client's profile.

**Total Ingress Flows**

The number of inflows coming to the client.

<b>Persistent Msgs In/sec</b>	The number of persistent incoming messages per second.
<b>Persistent Msgs Out/sec</b>	The number of persistent outgoing messages per second.
<b>Last Update</b>	The date and time of the last data update.
<b>Critical/Warning</b>	The number of critical alerts / warning alerts which also opens the <b>Alerts Table</b> .
<b>Non Persistent Msgs In/sec</b>	The number of non-persistent incoming messages per second.
<b>NonPersistent Msgs Out/sec</b>	The number of non-persistent outgoing messages per second.
<b>Uptime</b>	If the VPN's <b>Local Status</b> is <b>Up</b> , this field displays the length of time that the VPN has been up and running.
<b>Username</b>	The client's user name.
<b>Direct In Msgs /sec</b>	The number of non-persistent incoming messages per second.
<b>Direct Out Msgs /sec</b>	The number of non-persistent outgoing messages per second.

## Bridges

These displays provide process data for bridges configured on a VPN. Displays in this View are:

- **"Bridges Table"**: A tabular view of all available process performance data for all bridges configured on a VPN.
- **"Bridges - Diagram"**: Topological view of Solace network bridges that shows bridge broker connections and status.
- **"Bridge Summary"**: Current and historical metrics for a single bridge.

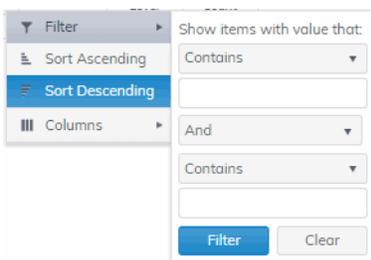
## Bridges Table

This display allows you to view data for all bridges configured for a VPN.

By default, a subset of available metrics is shown. Use **More Columns/Less Columns** to toggle to the complete set of metrics available (and back to the subset).

Select a broker and VPN from the drop-down menus. Use the check-boxes  to include / exclude **Enabled** and **Expired** bridges. Each table row is a different bridge.

Search by clicking the right side of a column heading/**Filter** to open the Search, Sort and Choose Columns dialog:



Rows listing bridges that are disabled or expired display with a shaded background. Double-click a row to drill down and investigate in the “[Bridge Summary](#)” display.

**Solace Bridges Table** 05-Mar-2019 09:01 No Alerts DATA

Broker: - All - VPN: blue-test Less Columns

Show Enabled Only:  Show: All Filter Bridge Name:  Bridges: 2

Broker	Local VPN	Bridge Name	Remote VPN	Remote R
VMR_114-US-EAST	blue-test	#bridge/v:solace/BridgeBroker1/15	BridgeBroker1	v:solace
VMR_114-US-EAST	blue-test	red_to_blue	red-test	v:ip-172-30-1-1

**Broker**

Displays the name of the broker

**Local VPN**

The name of the local VPN.

**Bridge Name**

The name of the bridge.

**Alert Level**

The current level of alerts in the row.

● Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.

● Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.

● Green indicates that no metrics have exceeded their alert thresholds.

**Alert Count**

The total number of active alerts for the process.

**Remote VPN**

The name of the remote VPN that is connected to the local VPN via the bridge.

**Remote Router**

The name of the remote broker.

**Admin State**

Indicates whether the bridge has been administratively enabled (via SolAdmin or the command line interface).

**Inbound Operational State**

The current inbound operational status of the bridge. (The administrator can turn off a bridge's input or output for maintenance or other reasons.)

**Outbound Operational State**

The current outbound operational status of the bridge. (The administrator can turn off a bridge's input or output for maintenance or other reasons.)

**Queue Operational State**

The current operational status of the queue.

<b>Connection Establisher</b>	Indicates whether the administrator created and configured the bridge directly on the broker using SolAdmin or the command line interface, or indirectly from another broker.
<b>Redundancy</b>	Displays whether the bridge is the <b>primary</b> bridge, the <b>backup</b> bridge, the <b>static</b> bridge (default bridge used when no other bridge is available), or whether it is the only bridge available ( <b>none</b> ).
<b>Uptime</b>	The current amount of time in which the bridge has been up and running.
<b>Client Name</b>	The name of the client.
<b>Connected Via Addr</b>	The local IP address and port used for the bridge.
<b>Connected Via Interface</b>	The name of the network interface used for the bridge.
<b>Client Direct Bytes Rcvd</b>	The number of bytes contained within direct messages received by the client via the bridge.
<b>Client Direct Bytes/sec Rcvd</b>	The number of bytes contained within direct messages received per second by the client via the bridge.
<b>Client Direct Bytes Sent</b>	The number of bytes contained within direct messages sent by the client via the bridge.
<b>Client Direct Bytes/sec Sent</b>	The number of bytes contained within direct messages sent per second by the client via the bridge.
<b>Client Direct Msgs/sec Rcvd</b>	The number of bytes contained within direct messages received per second by the client via the bridge.
<b>Client Direct Msgs Sent</b>	The number of direct messages sent by the client via the bridge.
<b>Client Direct Msgs/sec Sent</b>	The number of direct messages sent per second by the client via the bridge.
<b>Client NonPersistent Bytes Rcvd</b>	The number of bytes contained within non-persistent messages received by the client via the bridge.
<b>Client NonPersistent Bytes/sec Rcvd</b>	The number of bytes contained within non-persistent messages received per second by the client via the bridge.
<b>Client NonPersistent Bytes Sent</b>	The number of bytes contained within non-persistent messages sent by the client via the bridge.
<b>Client NonPersistent Bytes/sec Sent</b>	The number of bytes contained within non-persistent messages sent per second by the client via the bridge.
<b>Client NonPersistent Msgs Rcvd</b>	The number of non-persistent messages received by the client via the bridge.
<b>Client NonPersistent Msgs/sec Rcvd</b>	The number of non-persistent messages received per second by the client via the bridge.
<b>Client NonPersistent Msgs Sent</b>	The number of non-persistent messages sent by the client via the bridge.
<b>Client NonPersistent Msgs/sec Sent</b>	The number of non-persistent messages sent per second by the client via the bridge.
<b>Client Persistent Bytes Rcvd</b>	The number of bytes contained within persistent messages received by the client via the bridge.
<b>Client Persistent Bytes/sec Rcvd</b>	The number of bytes contained within persistent messages received per second by the client via the bridge.
<b>Client Persistent Bytes Sent</b>	The number of bytes contained within persistent messages sent by the client via the bridge.

<b>Client Persistent Bytes/sec Sent</b>	The number of bytes contained within persistent messages sent per second by the client via the bridge.
<b>Client Persistent Msgs Rcvd</b>	The number of persistent messages received by the client via the bridge.
<b>Client Persistent Msgs /sec Rcvd</b>	The number of persistent messages received per second by the client via the bridge.
<b>Client Persistent Msgs Sent</b>	The number of persistent messages sent by the client via the bridge.
<b>Client Persistent Msgs/sec Sent</b>	The number of persistent messages sent per second by the client via the bridge.
<b>Total Client Bytes Rcvd</b>	The number of bytes contained within all messages received by the client via the bridge.
<b>Total Client Bytes/sec Rcvd</b>	The number of bytes contained within all messages received per second by the client via the bridge.
<b>Total Client Bytes Sent</b>	The number of bytes contained within all messages sent by the client via the bridge.
<b>Total Client Bytes/sec Sent</b>	The number of bytes contained within all messages sent per second by the client via the bridge.
<b>Total Client Msgs Rcvd</b>	The total number of all messages received by the client via the bridge.
<b>Total Client Msgs/sec Rcvd</b>	The total number of all messages received per second by the client via the bridge.
<b>Total Client Msgs Sent</b>	The total number of all messages sent by the client via the bridge.
<b>Total Client Msgs/sec Sent</b>	The total number of all messages sent per second by the client via the bridge.
<b>Total Out Discards</b>	The total number of discarded outgoing messages sent by the client via the bridge.
<b>Total Out Discards/sec</b>	The total number of discarded outgoing messages sent per second by the client via the bridge.
<b>Total In Discards</b>	The total number of discarded incoming messages received by the client via the bridge.
<b>Total In Discards/sec</b>	The total number of discarded incoming messages received per second by the client via the bridge.

**Expired**

When checked, performance data about the bridge has not been received within the time specified (in seconds) in the `$solRowExpirationTime` field in the `conf\rtvapm_solmon.properties` file. The `$solRowExpirationTimeForDelete` field allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response from the bridge. To view/edit the current values, modify the following lines in the `.properties` file:

```
# Metrics data are considered expired after this
number of seconds
#
collector.sl.rtvapm.sub=$solRowExpirationTime:4
5
collector.sl.rtvapm.sub=$solRowExpirationTimeFo
rDelete:3600
```

In the example above, the **Expired** check box would be checked after 45 seconds, and the row would be removed from the table after 3600 seconds.

**Timestamp**

The date and time the row of data was last updated.

**Bridges - Diagram**

Use this topology view to monitor the health of your network bridges and VPNs. Quickly identify bridge and VPN connections, their health status and which resources their performance impacts. Drag and drop objects to arrange them on the screen (doing so does not logically impact the network bridges and VPNs). Arrows show the connections between VPNs and bridges.

Each object is a network bridge or VPN. Each is labeled with their name and color coded as follows:

- Red indicates that the object has one or more alerts in a critical state.
- Yellow indicates that the object has one or more alerts in a warning state.
- Green indicates that there are no alerts on the object.
- Gray indicates that the object is off-line.

**Save:** Saves the arrangement of the objects.

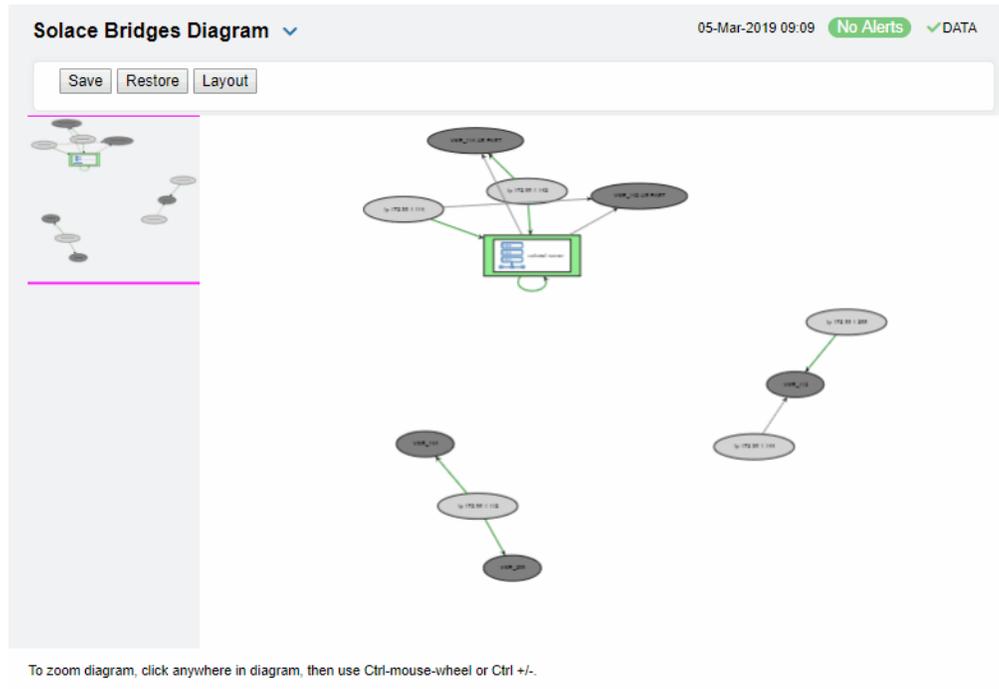
**Restore:** Returns objects to their previous positions.

**Layout:** Toggles between two types of layouts. One layout positions objects to the right so you might scroll in that direction to see them. The other layout pulls all the objects close together to the left, vertically, in hierarchical order.

Look at the miniature view in (upper left) to see all objects in either layout. Or zoom into the display using **Ctrl+/-** or **Ctrl+** mouse wheel.

Drill down to investigate in the ["Bridges Table"](#).

To monitor network brokers, VMRs and servers, see the [“Neighbors Diagram”](#).



## Bridge Summary

View current and historical performance and utilization metrics for a particular bridge on a VPN.

Select a broker, a VPN, and a bridge from the drop-down menus. Metric cards at the top of the displays show **Inbound and Outbound Client Messages per second**, **Ingress** and **Egress Discarded Messages**, and **Ingress** and **Egress KBs per second**.

You can hover over the metric cards to see more performance metrics and also drill down to see even more detail by clicking on them.

The trend graph traces current and historical performance metrics for the selected broker. The trend graph traces the performance metric you select: **Message Flow** or **Throughput**.

You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.



**Inbound Client Msgs/s**

The number of client messages received per second.

**Outbound Client Msgs/s**

The number of client messages sent per second.

**Ingress Discarded Client Msgs/s**

The number of discarded ingress messages per second.

**Egress Discarded Msgs/s**

The number of discarded egress messages per second.

**Inbound Client KB/s**

The amount of incoming client data, in KB per second.

**Outbound Client KB/s**

The amount of outgoing client data, in KB per second.

**Messages Flow Trend Graphs**

Traces the sum for the selected client.

- **Inbound Client Msgs/s:** The number of client messages received per second.
- **Outbound Client Msgs/s:** The number of client messages sent per second.

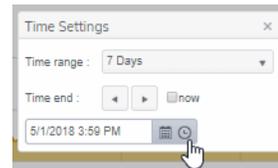
**Log Scale**

Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.

**Time Settings**

By default, the time range end point is the current time. To change the time range, click the **Time Settings**  and either:

- choose a **Time range** from 5 Minutes to 7 Days in the drop-down menu.
- specify begin/end dates using the calendar .
- specify begin/end time using the clock .



Toggle forward/backward in the trend graph per the period you choose (from the **Time range** drop-down menu) using arrows



Restore settings to current time by selecting **now** .

**Remote VPN**

The name of the remote VPN that is connected to the local VPN via the bridge.

**Expired**

When true, performance data about the bridge has not been received within the time specified (in seconds) in the **\$solRowExpirationTime** field in the **conf\rtvapm\_solmon.properties** file. The **\$solRowExpirationTimeForDelete** field allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response from the bridge. To view/edit the current values, modify the following lines in the **.properties** file:

```
# Metrics data are considered expired after this number
of seconds
#
collector.sl.rtvview.sub=$solRowExpirationTime:45
collector.sl.rtvview.sub=$solRowExpirationTimeForDelete:3
600
```

In the example above, the **Expired** check box would be checked after 45 seconds, and the row would be removed from the table after 3600 seconds.

**Address**

The IP address.

**Interface**

The interface ID.

**Queue Operational State**

Refer to Solace documentation for more information.

**Last Update**

The date and time of the last data update.

**Critical/Warning**

The number of critical alerts / warning alerts which also opens the **Alerts Table**.

**Remote Router**

The remote broker.

<b>Conn Establisher</b>	Refer to Solace documentation for more information.
<b>Inbound Operational State</b>	The current inbound operational status of the bridge. (The administrator can turn off a bridge's input or output for maintenance or other reasons.)
<b>Admin State</b>	Indicates whether the bridge has been administratively enabled (via SolAdmin or the command line interface).
<b>Client Name</b>	The name of the client.
<b>Redundancy</b>	Indicates whether the bridge is the <b>primary</b> bridge, the <b>backup</b> bridge, the <b>static</b> bridge (default bridge used when no other bridge is available), or whether it is the only bridge available ( <b>none</b> ).
<b>Outbound Op State</b>	The current outbound operational status of the bridge. (The administrator can turn off a bridge's input or output for maintenance or other reasons.)

## Endpoints

These displays list data for one or more endpoints configured on a VPN. Displays in this View are:

- ["Endpoints Table"](#)
- ["Endpoint Summary"](#)

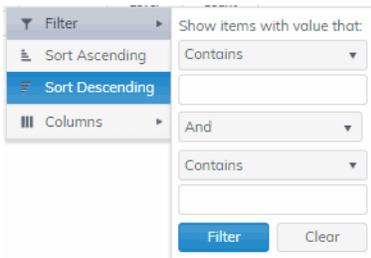
### Endpoints Table

View all endpoints configured on a VPN. Each row in the table lists the details for a specific endpoint.

By default, a subset of available metrics is shown. Use **More Columns/Less Columns** to toggle to the complete set of metrics available (and back to the subset).

Select a broker and VPN from the drop-down menus. Filter the table using the **Show Ingress Config Status Down Only** check-box  and use the **Show** drop-down menus to include **All**, **Expired** or **Unexpired**.

Search by clicking the right side of a column heading/**Filter** to open the Search, Sort and Choose Columns dialog:



You can click a column header to sort column data in numerical or alphabetical order, or double-click a row to drill down and investigate in the “Endpoint Summary” display.

**Solace Endpoints Table** 05-Mar-2019 09:13 No Alerts DATA

Broker:  VPN:  [Less Columns](#)

Show Ingress Config Status Down Only:  Show:  Filter Endpoint Name:  Endpoints: 2

Broker	VPN	Endpoint Name	Alert Level	Alert Count	Bind Count	Spooled Messages	Cur Spool Usage MB	High-Wa Mark M
solaceLoaner	Broker1	bridgeq	✔		1	2	0.0	
solaceLoaner	Broker1	q1	✔		0	0	0.0	

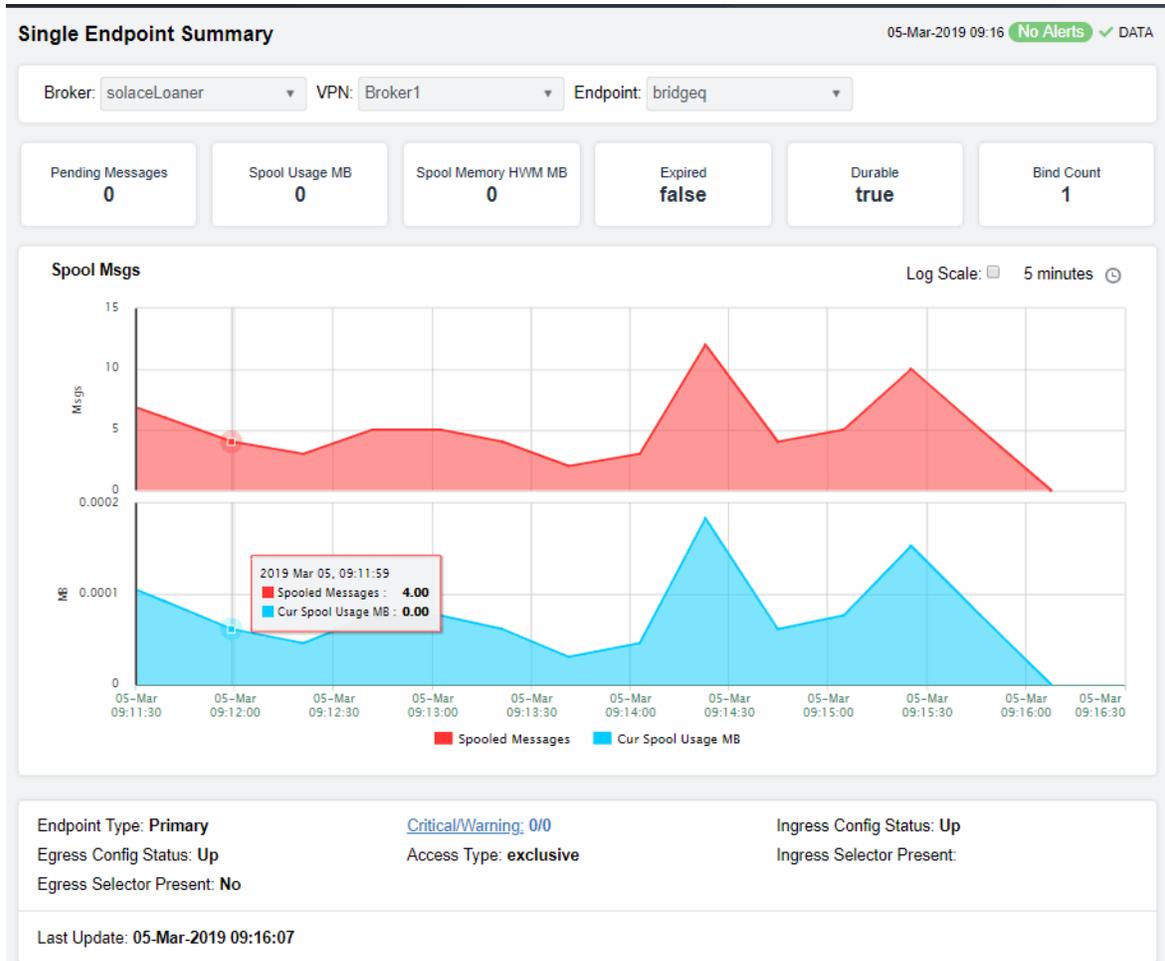
<b>Broker</b>	Displays the name of the broker
<b>VPN</b>	The name of the VPN.
<b>Endpoint Name</b>	The name of the endpoint.
<b>Alert Level</b>	The current alert severity in the row. <span style="color: red;">●</span> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold. <span style="color: yellow;">●</span> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold. <span style="color: green;">●</span> Green indicates that no metrics have exceeded their alert thresholds.
<b>Alert Count</b>	The total number of active alerts for the endpoint.
<b>Bind Count</b>	The total number of binds connected to the endpoint.
<b>Endpoint Type</b>	The type of endpoint (either queue or topic).
<b>Durable</b>	Displays whether or not the endpoint is durable (checked) or non-durable (unchecked). Durable endpoints remain after an broker restart and are automatically restored as part of an broker’s backup and restoration process.
<b>In Config Status</b>	Refer to Solace documentation for more information.
<b>Out Config Status</b>	Refer to Solace documentation for more information.
<b>Type</b>	Refer to Solace documentation for more information.
<b>Access Type</b>	Refer to Solace documentation for more information.

<b>Pending Messages</b>	The total number of pending messages on the endpoint.
<b>Spool Usage (MB)</b>	The total spool usage consumed on the endpoint (in megabytes).
<b>High Water Mark (MB)</b>	The highest level of spool usage on the endpoint (in megabytes).
<b>In Selector</b>	Refer to Solace documentation for more information.
<b>Out Selector</b>	Refer to Solace documentation for more information.
<b>Expired</b>	<p>When checked, performance data about the endpoint has not been received within the time specified (in seconds) in the <b>\$solRowExpirationTime</b> field in the <b>conf\rtvapm_solmon.properties</b> file. The <b>\$solRowExpirationTimeForDelete</b> field allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response from the endpoint. To view/edit the current values, modify the following lines in the <b>.properties</b> file:</p> <pre># Metrics data are considered expired after this number of seconds # collector.sl.rtvapm.sub=\$solRowExpirationTime:45 collector.sl.rtvapm.sub=\$solRowExpirationTimeForDelete:3600</pre> <p>In the example above, the <b>Expired</b> check box would be checked after 45 seconds, and the row would be removed from the table after 3600 seconds.</p>
<b>Time Stamp</b>	The date and time the row of data was last updated.

## Endpoint Summary

This display allows you to view endpoint information, message data, and a trend graph for pending and spool messages for a specific endpoint configured on a VPN. Choose a broker, a VPN, and an endpoint from the drop-down menus, and use the **Time Settings** to “zoom-in” or “zoom-out” on a specific time frame in the trend graph.

This display is provided by default and should be used if you do not want to collect message spool data for specific VPNs. However, if you do want to configure message spool monitoring for specific VPNs, then you should use the **Single Endpoint Summary Rates** display instead, which is not included in the navigation tree by default.

**Pending Messages**

The total number of pending messages on the endpoint.

**Spool Usage (MB)**

The current spool usage consumed on the endpoint (in megabytes).

**Spool Memory HWM MB**

Refer to Solace documentation for more information

**Expired**

When **true**, performance data about the endpoint has not been received within the time specified (in seconds) in the **\$solRowExpirationTime** field in the **conf\rtvapm\_solmon.properties** file. The **\$solRowExpirationTimeForDelete** field allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response from the endpoint. To view/edit the current values, modify the following lines in the **.properties** file:

```
# Metrics data are considered expired after this number
of seconds
#
collector.sl.rtvapm.sub=$solRowExpirationTime:45
collector.sl.rtvapm.sub=$solRowExpirationTimeForDelete
:3600
```

In the example above, the **Expired** check box would be checked after 45 seconds, and the row would be removed from the table after 3600 seconds.

**Durable**

Displays whether or not the endpoint is durable (checked) or non-durable (unchecked). Durable endpoints remain after an broker restart and are automatically restored as part of an broker's backup and restoration process.

**Bind Count**

The total number of binds connected to the endpoint.

**Trend Graphs**

Traces the sum of metrics for the endpoint.

- **Spooled Msgs:** The amount of spooled messages, in megabytes.
- **Cur Spool Usage:** The amount of space used by spooled messages, in megabytes.

**Log Scale**

Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.

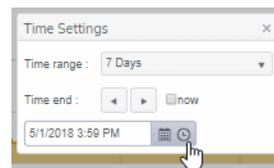
**Base at Zero**

Select to use zero (**0**) as the Y axis minimum for all graph traces.

**Time Settings**

By default, the time range end point is the current time. To change the time range, click the **Time Settings**  and either:

- choose a **Time range** from 5 Minutes to 7 Days in the drop-down menu.
- specify begin/end dates using the calendar .
- specify begin/end time using the clock .



Toggle forward/backward in the trend graph per the period you choose (from the **Time range** drop-down menu) using arrows  .

Restore settings to current time by selecting **now** .

<b>Endpoint Type</b>	The type of endpoint.
<b>Egress Config Status</b>	Refer to Solace documentation for more information.
<b>Egress Selector Present</b>	Refer to Solace documentation for more information.
<b>Last Update</b>	The date and time of the last data update.
<b>Critical/Warning</b>	The number of critical alerts / warning alerts which also opens the <b>Alerts Table</b> .
<b>Access Type</b>	Refer to Solace documentation for more information.
<b>Ingress Config Status</b>	Refer to Solace documentation for more information.
<b>Ingress Selector Present</b>	Refer to Solace documentation for more information.

## Capacity

These displays provide current broker capacity metrics, alert count and severity at the broker level. Displays in this View are:

- **"Capacity Table"**: View client, spool usage, incoming messages, outgoing messages, incoming bytes, and outgoing bytes data for all brokers.
- **"Capacity - Summary"**: View client, spool usage, incoming messages, outgoing messages, incoming bytes, and outgoing bytes data for a specific broker.
- **"Capacity Trends"**: View the broker capacity data for a specific broker in a trend graph format.

## Capacity Table

View current and HWM (high water mark for the last 30 days) capacity utilization data for all brokers.

By default, a subset of available metrics is shown. Use **More Columns/Less Columns** to toggle to the complete set of metrics available (and back to the subset).

You can view client, spool usage, incoming message, outgoing message, incoming bytes, and outgoing bytes data for the broker. Each table row is a different broker.

Search by clicking the right side of a column heading/**Filter** to open the Search, Sort and Choose Columns dialog:

The screenshot shows a dialog box for filtering data. On the left, there is a sidebar with four items: 'Filter' (selected), 'Sort Ascending', 'Sort Descending', and 'Columns'. The main area of the dialog is titled 'Show items with value that:' and contains two 'Contains' dropdown menus, one above the other, with an 'And' dropdown menu between them. At the bottom of the dialog, there are two buttons: 'Filter' and 'Clear'.

Double-click a row to drill down and investigate in the "Capacity - Summary" display.

**Broker Capacity Table** 05-Mar-2019 09:19 No Alerts ✓ DATA

Show: All Brokers: 1 Less Columns

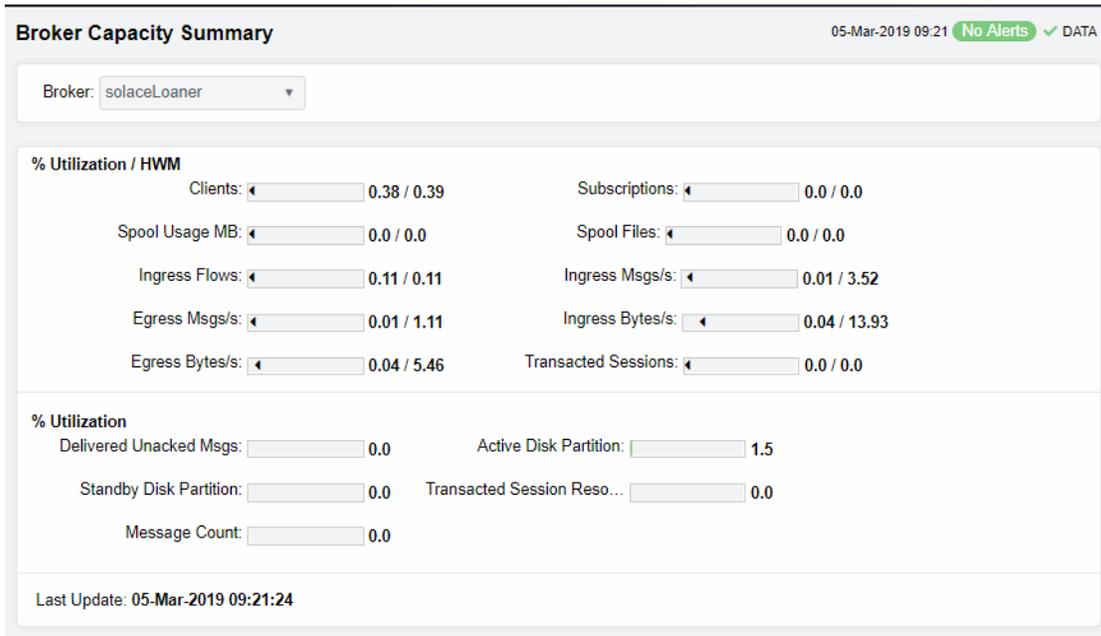
Broker	Alert Level	Alert Count	Current Client Connections	Connections HWM	Connections Max	Connections Reserved	Connections Used %
solaceLoaner	✓		35	35	9,000	135,000	

<b>Broker</b>	The name of the broker.
<b>Alert Level</b>	The maximum level of alerts in the row: <span style="color: red;">●</span> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold. <span style="color: yellow;">●</span> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold. <span style="color: green;">●</span> Green indicates that no metrics have exceeded their alert thresholds.
<b>Alert Count</b>	The total number of active alerts.
<b>Current Client Connections</b>	The current number of clients connected.
<b>Connections HWM</b>	The greatest number of connections in the last 30 days.
<b>Connections Max</b>	The greatest number of connections since the broker last started.
<b>Connections Reserved</b>	The current number of reserved connections.
<b>Connections Used %</b>	The current amount of connections used, in percent.
<b>Connections Used HWM %</b>	The greatest amount of connections used, in percent, in the last 30 days.
<b>Cur Spool Usage MB</b>	The current amount of used spool disk, in megabytes.
<b>Cur Spool Usage HWM</b>	The greatest amount of spool disk used in the last 30 days.
<b>Spool Disk Allocated</b>	The amount of allocated spool disk.
<b>Spool Reserved</b>	The amount of reserved spool disk.
<b>Current Spool Usage %</b>	The current amount of used spool disk, in percent.
<b>Current Spool Usage % HWM</b>	The greatest amount of used spool disk in the last 30 days, in percent.

<b>Delivered Unacked Msgs Util %</b>	Refer to Solace documentation for more information.
<b>Ingress Flow Count</b>	The number of ingress flows.
<b>Ingress Flow HWM</b>	The greatest number of ingress flows in the last 30 days.
<b>Ingress Flows Allowed</b>	The maximum number of ingress flows allowed.
<b>Ingress Flow Count %</b>	The amount of ingress flows in percent.
<b>Ingress Flow Count HWM %</b>	The greatest amount of ingress flows in the last 30 days, in percent.
<b>Ingress Msgs/s</b>	The number of ingress messages per second.
<b>Ingress Msgs/s HWM</b>	The greatest number of ingress messages per second in the last 30 days.
<b>Max Ingress Msgs/s</b>	The maximum number of ingress flows per second allowed.
<b>Ingress Msgs %</b>	The amount of ingress messages in percent.
<b>Ingress Msgs/s HWM %</b>	The greatest amount of ingress messages in the last 30 days, in percent.
<b>Cur Egress Msgs/s</b>	The number of egress messages per second.
<b>Egress Msgs/s HWM</b>	The greatest number of egress messages per second in the last 30 days.
<b>Max Egress Msgs/s</b>	The maximum number of egress flows per second allowed.
<b>Egress Msgs %</b>	The amount of egress messages in percent.
<b>Egress Msgs/s HWM %</b>	The greatest amount of ingress messages in the last 30 days, in percent.
<b>Cur Egress Bytes/s</b>	The amount of egress in bytes per second.
<b>Egress Bytes/s HWM</b>	The greatest amount of egress, in bytes per second, in the last 30 days, in percent.
<b>Expired</b>	<p>When checked, performance data about the VPN has not been received within the time specified (in seconds) in the <b>\$solRowExpirationTime</b> field in the <b>conf\rtvapm_solmon.properties</b> file. The <b>\$solRowExpirationTimeForDelete</b> field allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response from the VPN. To view/edit the current values, modify the following lines in the <b>.properties</b> file:</p> <pre># Metrics data are considered expired after this number of seconds # collector.sl.rtvview.sub=\$solRowExpirationTime:45 collector.sl.rtvview.sub=\$solRowExpirationTimeForDelete:3600</pre> <p>In the example above, the <b>Expired</b> check box would be checked after 45 seconds, and the row would be removed from the table after 3600 seconds.</p>
<b>Time Stamp</b>	The date and time the row of data was last updated.

## Capacity - Summary

This display, a pivoted view of the “Capacity Table”, allows you to view current and HWM (high water mark for the last 30 days) capacity utilization data for a single broker. Select a broker from the drop-down menu to view client, spool usage, incoming message, outgoing message, incoming bytes, and outgoing bytes data for the broker.



### % Utilization/HWM

These values show high water marks (peak capacity utilization) for the last 30 days.

#### Clients

The current number of clients connected to the broker.

#### Spool Files

The highest number of spool files on the broker in the past 30 days.

#### Egress Msgs/s

The highest number of outgoing messages per second on the broker in the past 30 days.

#### Transacted Sessions

The highest number of transacted sessions on the broker in the last 30 days.

#### Subscriptions

The highest number of subscriptions on the broker in the last 30 days.

#### Ingress Flows

The highest number of inflows on the broker in the last 30 days.

#### Ingress Bytes/s

The highest amount of inflows, in bytes per second, on the broker in the past 30 days.

#### Spool Usage MB

The highest amount of spool utilization, in megabytes per second, on the broker in the past 30 days.

#### Ingress Msgs/s

The highest number of incoming messages per second on the broker in the past 30 days.

#### Egress Bytes/s

The highest number of outgoing messages per second on the broker in the past 30 days.

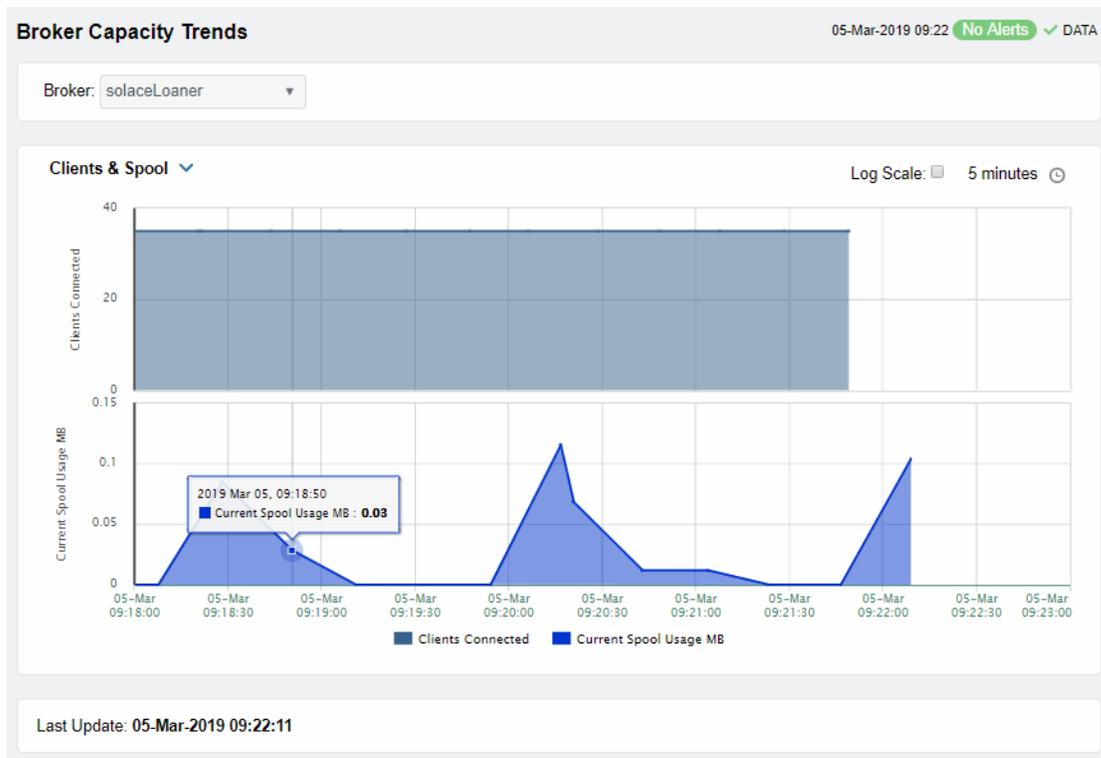
**% Utilization**

These values show current capacity utilization.

<b>Delivered Unacked Msgs</b>	The current number of delivered messages that were not acknowledged divided by the maximum number of delivered messages that were not acknowledged allowed on the broker.
<b>Transacted Sessions Reso...</b>	The current number of transacted sessions that were resolved on the broker.
<b>Active Disk Partition</b>	The percentage of available active disk partition that is used.
<b>Message Count</b>	The current number of messages on the broker.
<b>Standby Disk Partition</b>	The percentage of available standby disk partition that has been used.
<b>Last Update</b>	The date and time of the last data update.

**Capacity Trends**

This display allows you to view a trend graph that traces broker performance data for clients & spool data, message flow and throughput. Select a broker and a performance metric from the drop-down menus.



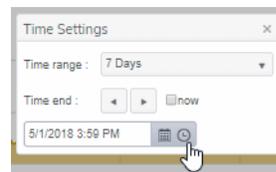
**Clients & Spool**

The trend graph traces the following performance metrics:

**Clients Connected:** The current number of clients connected to the broker.

**Current Spool Usage:** The current spool usage, in megabytes, on the broker.

- Message Flow** The trend graph traces the following:  
**Ingress Msgs/sec**: The number of incoming messages per second on the broker.  
**Egress Msgs/sec**: The number of outgoing messages per second on the broker.
- Throughput** The trend graph traces the following:  
**Ingress KB/sec**: The amount of incoming per second, in KB, on the broker.  
**Egress KB/sec**: The number of outgoing data per second, in KB, on the broker.
- Log Scale** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Base at Zero** Select to use zero (0) as the Y axis minimum for all graph traces.
- Time Settings** By default, the time range end point is the current time. To change the time range, click the **Time Settings**  and either:
- choose a **Time range** from 5 Minutes to 7 Days in the drop-down menu.
  - specify begin/end dates using the calendar  .
  - specify begin/end time using the clock  .



Toggle forward/backward in the trend graph per the period you choose (from the **Time range** drop-down menu) using arrows   .  
 Restore settings to current time by selecting **now**  .

## Syslog Events

The Solace Syslog Events Table allows you to supervise the last Syslog event messages from the Solace Message Brokers that have been configured for syslog monitoring. See the Solace product documentation for an in depth description of Syslog monitoring in Solace products and how to configure the Message Brokers and the Syslog destination.

This display requires the Solace Event Module from the RTView Solace Monitor to be properly configured with a Syslog destination and running. See ["Solace Event Module"](#) for additional information.

## Syslog Events Table

This display lists all Syslog events collected from all Solace brokers. Each row in the table is a different message. Use the drop-down menus to filter the list by **Connection**, **Scope** and alert **Severity** level. Filter messages per single broker or all brokers. Click a column header to sort column data in numerical, alphabetical or chronological order.

Search by clicking the right side of a column heading/**Filter** to open the Search, Sort and Choose Columns dialog:

Provisioning
Solace Syslog Events Table
27-Mar-2019 12:33 ✓ DATA

Connection: - All -

More Columns

Scope: - All -

Severity: - All -

Show: All

Events: 3,194

time_stamp	Connection	Solace Scope	Host	Facility	Solace Event...	Syslog Seve...	Type	Message	Additional Fl...	Cl
26-Mar-2019 12: VMR-47	SYSTEM	ip-172-30-1-144	19	SYSTEM_AUTH	notice	SYSLOG_MSG	Tue Mar 26 19:3	("sessionType":		
26-Mar-2019 12: VMR-118	SYSTEM	ip-172-30-1-112	19	SYSTEM_AUTH	notice	SYSLOG_MSG	Tue Mar 26 19:3	("sessionType":		
26-Mar-2019 12: VMR-47	SYSTEM	ip-172-30-1-144	19	SYSTEM_AUTH	notice	SYSLOG_MSG	Tue Mar 26 19:3	("sessionType":		
26-Mar-2019 12: VMR-118	SYSTEM	ip-172-30-1-112	19	SYSTEM_AUTH	notice	SYSLOG_MSG	Tue Mar 26 19:3	("sessionType":		
26-Mar-2019 12: VMR-47	SYSTEM	ip-172-30-1-144	19	SYSTEM_AUTH	notice	SYSLOG_MSG	Tue Mar 26 19:3	("sessionType":		
26-Mar-2019 12: VMR-118	SYSTEM	ip-172-30-1-112	19	SYSTEM_AUTH	notice	SYSLOG_MSG	Tue Mar 26 19:3	("sessionType":		
26-Mar-2019 12: VMR-47	SYSTEM	ip-172-30-1-144	19	SYSTEM_AUTH	notice	SYSLOG_MSG	Tue Mar 26 19:3	("sessionType":		
26-Mar-2019 12: VMR-118	SYSTEM	ip-172-30-1-112	19	SYSTEM_AUTH	notice	SYSLOG_MSG	Tue Mar 26 19:3	("sessionType":		
26-Mar-2019 12: VMR-47	SYSTEM	ip-172-30-1-144	19	SYSTEM_AUTH	notice	SYSLOG_MSG	Tue Mar 26 19:3	("sessionType":		
26-Mar-2019 12: VMR-118	SYSTEM	ip-172-30-1-112	19	SYSTEM_AUTH	notice	SYSLOG_MSG	Tue Mar 26 19:3	("sessionType":		
26-Mar-2019 12: VMR-47	SYSTEM	ip-172-30-1-144	19	SYSTEM_AUTH	notice	SYSLOG_MSG	Tue Mar 26 19:3	("sessionType":		
26-Mar-2019 12: VMR-118	SYSTEM	ip-172-30-1-112	19	SYSTEM_AUTH	notice	SYSLOG_MSG	Tue Mar 26 19:3	("sessionType":		
26-Mar-2019 12: VMR-47	SYSTEM	ip-172-30-1-144	19	SYSTEM_AUTH	notice	SYSLOG_MSG	Tue Mar 26 19:3	("sessionType":		
26-Mar-2019 12: VMR-118	SYSTEM	ip-172-30-1-112	19	SYSTEM_AUTH	notice	SYSLOG_MSG	Tue Mar 26 19:3	("sessionType":		

Page 1 of 80

1 - 40 of 3194 items

### Connection

Select the connection string assigned when the message brokers connection properties were added with the RTView Configuration Application.

### More/Fewer Columns

Switches to another syslog events table display containing the full set of columns coming from Syslog.

<b>Scope:</b>	<p>This drop down selects the type of the event. The SYSTEM events are coming from conditions related to the state of the message broker. VPN events are events with the state of the message brokers VPNs. CLIENT events refer to the state of clients executions in the messaging infrastructure.</p> <p>Available options are:</p> <ul style="list-style-type: none"><li>• SYSTEM</li><li>• VPN</li><li>• CLIENT</li><li>• ALL shows messages from all sources.</li></ul>
<b>Severity:</b>	<p>Selects the severity level of the events that will be presented in the table. All options go from the less severe to the most important to the health of the systems unless one specifies one single type of severity. For instance, Warning will only show the events that are defined as Warning, filtering out events more damaging, whereas Warning or higher will show all Syslog events that are either Warning, Error, Alert or Emergency. To avoid missing any key event, selection of Warning or higher is recommended.</p> <p>Available options are:</p> <ul style="list-style-type: none"><li>• INFO</li><li>• NOTICE</li><li>• NOTICE or higher</li><li>• WARN</li><li>• WARN or higher</li><li>• ERROR</li><li>• ERROR or higher</li><li>• CRITICAL</li><li>• ALERT</li><li>• EMERGENCY</li><li>• ALL shows messages regardless of severity level from all sources.</li></ul>
<b>Show:</b>	<p>Selects the Expiration flag of the event. Due to the large number of events that can exist, it is recommended to select Unexpired Only to see exclusively the events that are active.</p> <p>Available options are:</p> <ul style="list-style-type: none"><li>• Expired Only</li><li>• Unexpired Only</li><li>• ALL shows expired and unexpired messages from all sources.</li></ul>
<b>Events:</b>	<p>The number of events currently shown in the table.</p>
<b>Time Stamp</b>	<p>The date and time the row of data was last updated.</p>





## CHAPTER 8 RTView DataServer for TIBCO

The RTView DataServer for TIBCO provides a way to create connections and modify default configuration settings for solution packages and sends collected data to RTViewCentral. RTViewCentral contains the displays associated with the RTView DataServer for TIBCO which you use to monitor your TIBCO components. Both the display server user interface and the HTML user interface are described here.

The RTView *DataCollector* for TIBCO is available for use with the RTView DataServer for TIBCO. RTView DataCollector for TIBCO is used for collecting data and sending it to one or more RTView DataServers. The RTView DataCollector for TIBCO is also useful if you need to distribute data collection.

For an overview and details about configuring RTView Enterprise, including RTViewCentral, RTView DataServers, RTView DataCollectors and solution packages, see the *RTView Enterprise Configuration Guide*.

RTViewCentral contains the following solution packages and associated displays that will be populated with data collected via the RTView DataServer for TIBCO:

- ["TIBCO ActiveMatrix"](#)
- ["TIBCO ActiveSpaces"](#)
- ["TIBCO ActiveSpaces - HTML"](#)
- ["TIBCO ActiveSpaces \(2.x\)"](#)
- ["TIBCO ActiveSpaces \(2.x\) - HTML"](#)
- ["TIBCO Adapters"](#)
- ["TIBCO Adapters - HTML"](#)
- ["TIBCO BusinessEvents"](#)
- ["TIBCO BusinessEvents - HTML"](#)
- ["TIBCO BusinessWorks"](#)
- ["TIBCO BusinessWorks - HTML"](#)
- ["TIBCO BusinessWorks 5 Monitor - HTML"](#)
- ["TIBCO Enterprise Message Service"](#)
- ["TIBCO Enterprise Message Service - HTML"](#)
- ["TIBCO FTL"](#)
- ["TIBCO FTL - HTML"](#)
- ["TIBCO Hawk"](#)
- ["TIBCO Hawk - HTML"](#)
- RTView Manager is also included. For details, see ["RTView Manager"](#).

**Note:** This document assumes familiarity with the products monitored. For additional details, refer to vendor documentation.

---

## TIBCO ActiveMatrix

The TIBCO ActiveMatrix Views can be found under **Components** tab > **Middleware**. The displays within the Views will be populated with data once the Solution Package for TIBCO ActiveMatrix is configured in the RTView DataServer for TIBCO and the RTView DataServer for TIBCO is connected to RTViewCentral.

---

## TIBCO ActiveSpaces

The following TIBCO ActiveSpaces Views can be found under **Components** tab > **Middleware** > **TIBCO ActiveSpaces**:

- **"All Realms Views"**: The displays in this View provide detailed data for all realms in a heatmap and tabular format.
- **"Single Realm Views"**: The displays in this View provide detailed current and historical data for a particular realm in tabular and trend graph format.
- **"All Node Views"**: The displays in this View provide detailed data for all nodes (in a specific realm) in a heatmap or tabular format.
- **"Single Node Views"**: The display in this View provides detailed data for a single node for a particular realm/node combination.
- **"All Proxy Views"**: The displays in this View provide detailed data for all proxies in a heatmap and tabular format.
- **"Single Proxy Views"**: The display in this View provides detailed data for a single proxy for a realm/proxy combination.
- **"All Keeper Views"**: The displays in this View provide detailed data for all keepers in a heatmap or tabular format.
- **"Single Keeper Views"**: The display in this View provides detailed data for a single keeper in a summary format.

### All Realms Views

These displays provide detailed data for all realms in a heatmap and tabular format. Displays in this View are:

- **"All Realms Heatmap"**: A heatmap view of all realms.
- **"All Realms Table"**: A tabular view of your realms and their associated metrics.

### All Realms Heatmap

This heatmap display provides an easy-to-view interface that allows you to quickly identify the current status of each of your realms for each available metric. You can view the realms in the heatmap based on the following metrics: current alert severity, alert count, CPU usage, memory usage, operations completed, and operations failed. By default, this display shows the heatmap based on the **Alert Severity** metric.

You can use the **Names** check-box  to include or exclude labels in the heatmap, and you can mouse over a rectangle to see additional metrics for a realm. Clicking one of the rectangles in the heatmap opens the “**Realm Summary**” display, which allows you to see additional details for the selected realm.



**Title Bar** (possible features are):

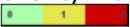
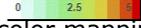
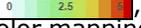
- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

### Fields and Data:

- Names** Select this check box to display the names of the realms at the top of each rectangle in the heatmap.
- Log** Select this check box to use a logarithmic scale, rather than a linear scale, to map from the selected metric value for a cell to the color for the cell. **Log** provides another way to distribute and differentiate values that you might not be able to see on a linear scale due to the dominant nature of large values in a linear scale.
- Auto** When checked, the values of the selected metric are auto-scaled to its highest defined value. When unchecked, the values of the selected metric display based on the threshold defined for the alert associated with the selected metric. Selecting **Auto** helps to visualize the range of the values currently present for the selected metric instead of the threshold of the alert that has been associated with the metric. All metrics that have not been associated in the heatmap defaults with alerts use a monochromatic color gradient bar (whites and greens). All metrics that have been associated in the heatmap defaults with alerts use a multi-chromatic color gradient bar (reds, yellows, white, and greens).
- Metric** Choose a metric to view in the display.

<b>Alert Severity</b>	<p>The current alert severity. Values range from <b>0</b> - <b>2</b>, as indicated in the color gradient  bar, where <b>2</b> is the highest Alert Severity:</p> <ul style="list-style-type: none"> <li> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Alert Count</b>	<p>The total number of alarm and warning alerts in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>CPU Usage</b>	<p>The milliseconds of CPU time accumulated by the process after each update interval. The color gradient bar  , populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of <b>TdgRealmServerCpuUsageHigh</b>. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>Memory</b>	<p>The amount of memory used in the realm. The color gradient bar  , populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of <b>TdgRealmServerMemoryUseHigh</b>. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>Ops Completed</b>	<p>The number of operations completed in the realm. The color gradient bar  , populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of operations in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>Ops Failed</b>	<p>The number of failed operations in the realm. The color gradient bar  , populated by the current heatmap, shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of operations that have failed in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>

## All Realms Table

The table in this display provides a view of all of your realms and their associated metric data including alert level, alert count, and the current value of each gathered metric. You can click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by clicking a row to view details for the selected realm in the ["Realm Summary"](#) display.

Realm	Alert Level	Alert Count	Expired	Grid Name	Server Cpu (%)	Server Memory (K...)	Server Id	Server Uptime	Server Version	Co
B44-Realm1			<input type="checkbox"/>	_default	1.9	297960	e59808d6-8a	938	6.1.0 V5	

**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu , Table open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the selected domain. Refer to TIBCO ActiveSpaces documentation for more information regarding these fields.

**All Realms Table**

- Realm** The name of the realm.
- Alert Level** The current alert severity.
  - Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
  - Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
  - Green indicates that no metrics have exceeded their alert thresholds.
- Alert Count** The total number of alerts for the host.
- Expired** When checked, performance data has not been received within the time specified (in seconds) in the **Expire Time** field in the **Duration** region in the RTView Configuration Application > (Project Name) > **Solution Package Configuration** > **TIBCO Active Spaces** > **DATA STORAGE** tab. The **Delete Time** field (also in the **Duration** region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.

<b>Grid Name</b>	The grid name associated with the realm.
<b>Server CPU (%)</b>	The percentage of milliseconds of CPU time accumulated by the server after the last update interval.*
<b>Server Memory (KB)</b>	The amount of memory used on the server, in kilobytes.
<b>Server ID</b>	The ID of the server.
<b>Server Uptime</b>	The amount of time since the server was last started.
<b>Server Version</b>	The version of the server.
<b>Ops Completed</b>	The total number of operations completed in the realm.*
<b>Ops Completed Rate</b>	The rate of operations completed in the realm.*
<b>Ops Failed</b>	The total number of operations failed in the realm.*
<b>Ops Failed Rate</b>	The rate of operations failed in the realm.*
<b>Node Live Data</b>	The size of the node's live data.*
<b>Txn Commits</b>	The number of transaction commits in the realm.*
<b>Txn Commits Rate</b>	The rate of transaction commits in the realm.*
<b>Txn Rollbacks</b>	The number of transaction rollbacks in the realm.*
<b>Txn Rollbacks Rate</b>	The rate of transaction rollbacks in the realm.*
<b>Timestamp</b>	The date and time the row data was last updated.

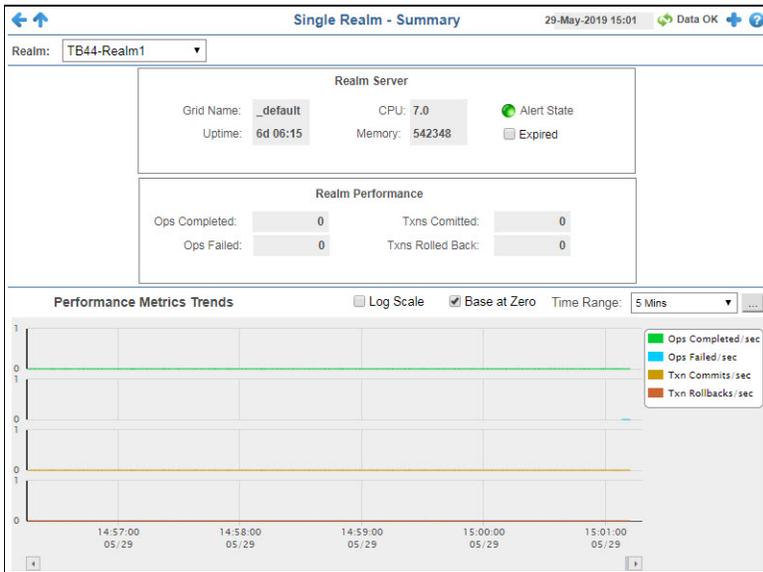
## Single Realm Views

These displays provide detailed current and historical data for a particular realm in tabular and trend graph format. Displays in this View are:

- **"Realm Summary"**: This display allows you to view metrics and trend data for a particular realm.
- **"Realm Servers"**: The tables in this display provide a view of your group server metrics, the group metrics, the persist server metrics, and the persist metrics for a selected realm.

### Realm Summary

This display provides a view of the current and historical metrics for a single realm. The trend graph in the bottom half of the display traces the rate of completed operations, the rate of failed operations, the rate of transactions that are committed, and the rate of transactions that are rolled back.



**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the selected domain. Refer to TIBCO ActiveSpaces documentation for more information regarding these fields.

**Filter By:**

The display might include these filtering options:

**Realm** Select the realm for which you want to show data in the display.

**Fields and Data:**

**Realm Server**

**Grid Name** The name of the grid.

**CPU** The CPU usage percentage.

<b>Alert State</b>	<p>The current alert severity.</p> <ul style="list-style-type: none"> <li><span style="color: red;">●</span> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li><span style="color: yellow;">●</span> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li><span style="color: green;">●</span> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Uptime</b>	The amount of time since the server was started.
<b>Memory</b>	The amount of memory used on the server.
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > (Project Name) > <b>Solution Package Configuration</b> > <b>TIBCO Active Spaces</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.

### Realm Performance

<b>Ops Completed</b>	The number of operations completed in the realm.*
<b>Op Failed</b>	The number of operation failed in the realm.*
<b>Txns Committed</b>	The total number of transactions committed in the realm.*
<b>Txns Rolled Back</b>	The number of transactions rolled back in the realm.*

### Trends

Traces the following:

**Ops Completed/sec**-- traces the number of operations completed per second.

**Ops Failed/sec**-- traces the number of operations failed per second.

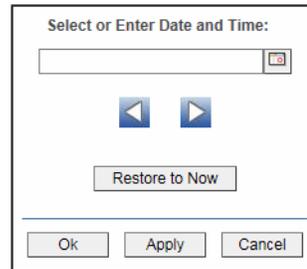
**Txn Commits/sec**-- traces the number of transactions committed per second.

**Txn Rollbacks/sec**-- traces the number of transactions rolled back per second.

**Log Scale** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.

**Base at Zero** Select to use zero (0) as the Y axis minimum for all graph traces.

**Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar  .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

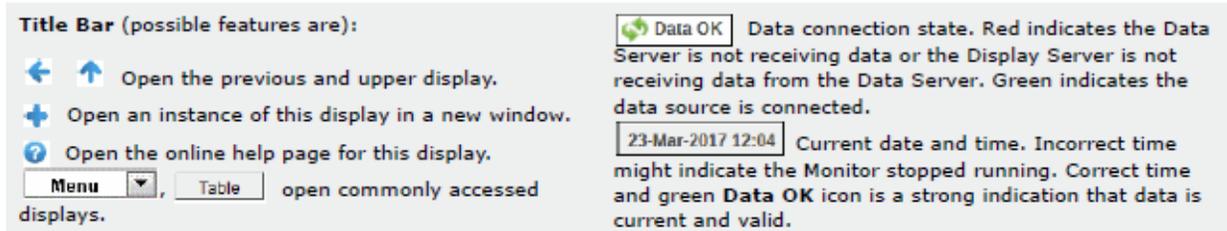
Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Realm Servers

The tables in this display provide a view of your group server metrics, the group metrics, the persist server metrics, and the persist metrics for a selected realm. You can click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by clicking a row to view additional details in the "Realm Summary" display.

Single Realm - Servers								
Realm: TB44-Realm1								
Group Server								
Connection	Group	PID	Host	CPU Utilization	Messages Rcvd Rate	Messages Sent Rate	time_stamp	
TB44-Realm1	1036	1532	TESTBED-44		0	0	11-Jun-2019 14:35:12	
Group Metrics								
time_stamp	Group	Metric	value					
11-Jun-2019 14:35:12	1036	FORMAT_UNAVAILABLE	0					
11-Jun-2019 14:35:12	1036	DATA_LOST	0					
11-Jun-2019 14:35:12	1036	BYTES_RECEIVED	0					
11-Jun-2019 14:35:12	1036	BYTES_SENT	0					
11-Jun-2019 14:35:12	1036	QUEUE_DISCARDS	0					
11-Jun-2019 14:35:12	1036	QUEUE_BACKLOG	0					
11-Jun-2019 14:35:12	1036	PROCESS_VM_KB	2651936					
11-Jun-2019 14:35:12	1036	PROCESS_PFAK_RSS_KB	421316					
Persist Server								
Connection	Persist	PID	Host	CPU Utilization	Messages Rcvd Rate	Messages Sent Rate	time_stamp	
TB44-Realm1	2965	0	TESTBED-44		0	0	11-Jun-2019 14:35:12	
Persist Metrics								
time_stamp	Persist	Metric	value					
11-Jun-2019 14:35:12	libdg_AC439694-761E-470F-9123-D8ADF508DE15	MESSAGE_SIZE	0					
11-Jun-2019 14:35:12	libdg_AC439694-761E-470F-9123-D8ADF508DE15	MESSAGE_COUNT	0					
11-Jun-2019 14:35:12	libdg_AC439694-761E-470F-9123-D8ADF508DE15	_StoreDispatcher Event Q	0					
11-Jun-2019 14:35:12	libdg_AC439694-761E-470F-9123-D8ADF508DE15	_StoreDispatcher Event Q	0					
11-Jun-2019 14:35:12	libdg_AC439694-761E-470F-9123-D8ADF508DE15	DYNAMIC_FORMATS	85					
11-Jun-2019 14:35:12	libdg_AC439694-761E-470F-9123-D8ADF508DE15	USER_CPU_TIME	83350000					
11-Jun-2019 14:35:12	libdg_AC439694-761E-470F-9123-D8ADF508DE15	SYSTEM_CPU_TIME	248080000					
11-Jun-2019 14:35:12	libdg_AC439694-761E-470F-9123-D8ADF508DE15		33420000					




---

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the selected domain. Refer to TIBCO ActiveSpaces documentation for more information regarding these fields.

---

### Filter By:

**Realm** Select the realm for which you want to view data.

### Group Server Table:

**Connection** The name of the connection

**Group** The name/ID of the group.

**PID** The process ID of the group server.\*

**Host** The name of the host.\*

**CPU Utilization** The CPU utilization for the group server.\*

**Messages Rcvd Rate** The rate of messages received on the group server.\*

**Messages Sent Rate** The rate of messages sent on the group server.\*

**time\_stamp** The date and time the row data was last updated.

### Group Metrics Table:

**time\_stamp** The date and time the row data was last updated.

**Group** The name/ID of the group.\*

**Metric** The name of the metric.\*

**value** The value of the metric.\*

### Persist Server Table:

**Connection** The name of the connection.\*

**Persist** The name of the persist process.\*

**PID** The process ID of the persist process.\*

**Host** The name of the host.\*

**CPU Utilization** The CPU utilization percentage on the persist server.\*

<b>Messages Rcvd Rate</b>	The rate of messages received on the persist server.*
<b>Messages Sent Rate</b>	The rate messages sent on the persist server.*
<b>time_stamp</b>	The date and time the row data was last updated.

**Persist Metrics Table:**

<b>time_stamp</b>	The date and time the row data was last updated.
<b>Persist</b>	The persist process name.*
<b>Metric</b>	The name of the metric.*
<b>Value</b>	The value of the metric.*

## All Node Views

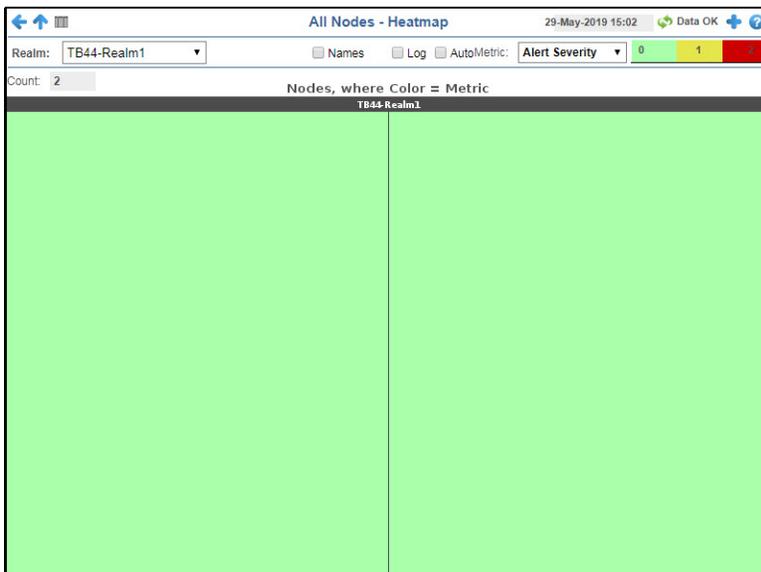
These displays provide detailed data for all nodes (in a specific realm) in a heatmap or tabular format. Displays in this View are:

- **"All Nodes Heatmap"**: A heatmap view of all nodes contained within a particular realm.
- **"All Nodes Table"**: A tabular view of all nodes (contained within a particular realm) and their associated metrics.

### All Nodes Heatmap

This heatmap display provides an easy-to-view interface that allows you to quickly identify the current status of each of your nodes for each available metric. You can view the nodes in the heatmap based on the following metrics: current alert severity, alert count, CPU usage, memory usage, rate of failed operations, rate of completed operations, and rate of transactions that were rolled back. By default, this display shows the heatmap based on the **Alert Severity** metric.

You can use the **Names** check-box  to include or exclude labels in the heatmap, and you can mouse over a rectangle to see additional metrics for a node. Clicking one of the rectangles in the heatmap opens the **"Node Summary"** display, which allows you to see additional details for the selected node.



#### Title Bar (possible features are):

-   Open the previous and upper display.
-  Open an instance of this display in a new window.
-  Open the online help page for this display.
- open commonly accessed displays.

 **Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

#### Filter By:

**Realm** Select the realm for which you want to see data.

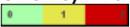
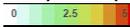
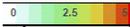
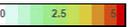
#### Fields and Data:

**Names** Select this check box to display the names of the realms at the top of each rectangle in the heatmap.

**Log** Select this check box to use a logarithmic scale, rather than a linear scale, to map from the selected metric value for a cell to the color for the cell. **Log** provides another way to distribute and differentiate values that you might not be able to see on a linear scale due to the dominant nature of large values in a linear scale.

**Auto** When checked, the values of the selected metric are auto-scaled to its highest defined value. When unchecked, the values of the selected metric display based on the threshold defined for the alert associated with the selected metric. Selecting Auto helps to visualize the range of the values currently present for the selected metric instead of the threshold of the alert that has been associated with the metric. All metrics that have not been associated in the heatmap defaults with alerts use a monochromatic color gradient bar (whites and greens). All metrics that have been associated in the heatmap defaults with alerts use a multi-chromatic color gradient bar (reds, yellows, white, and greens).

**Metric** Choose a metric to view in the display.

<b>Alert Severity</b>	<p>The current alert severity. Values range from <b>0</b> - <b>2</b>, as indicated in the color gradient  bar, where <b>2</b> is the highest Alert Severity:</p> <ul style="list-style-type: none"> <li> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Alert Count</b>	<p>The total number of alarm and warning alerts in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>CPU Usage</b>	<p>The milliseconds of CPU time accumulated by the process after the last update interval. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of <b>TdgNodeCpuUsageHigh</b>. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>Memory</b>	<p>The memory usage for the node. The color gradient bar , populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of <b>TdgNodeMemoryUseHigh</b>. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>Failed Op Rate</b>	<p>The rate of failed operations. The color gradient bar , populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of <b>TdgNodeOpsFailedRateHigh</b>. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>Completed Op Rate</b>	<p>The rate of completed operations. The color gradient bar , populated by the current heatmap, shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of <b>TdgNodeOpsCompletedRateLow</b>. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>Txn Rollback Rate</b>	<p>The rate of transactions that have been rolled back. The color gradient bar , populated by the current heatmap, shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of <b>TdgNodeTxnRollbackRateHigh</b>. The middle value in the gradient bar indicates the middle value of the range.</p>

## All Nodes Table

The table in this display provides a view of all nodes and their associated metric data in a specific realm. You can click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by double-clicking a row to view details for the selected node in the "Node Summary" display.

Realm	Node	Alert Level	Alert Count	Expired	Client Status	Client Id	Host	Copyset	PID	Node Ready
TB44-Realm1n1				<input type="checkbox"/>	RUNNING	1042	TESTBED-44 set1		13567	<input checked="" type="checkbox"/>
TB44-Realm1n2				<input type="checkbox"/>	RUNNING	1044	TESTBED-44 set2		13568	<input checked="" type="checkbox"/>

**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the selected domain. Refer to TIBCO ActiveSpaces documentation for more information regarding these fields.

**Filter By:**

The display might include these filtering options:

- Realm** Select the realm (containing the node) for which you want to show data in the display.
- Count** The number of realms found in the search and listed in the table.

**All Nodes Table**

- Realm** The name of the realm.
- Node** The name of the node.
- Alert Level** The current alert severity.
  - Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
  - Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
  - Green indicates that no metrics have exceeded their alert thresholds.

<b>Alert Count</b>	The ID of the node.*
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > (Project Name) > <b>Solution Package Configuration</b> > <b>TIBCO Active Spaces</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Client Status</b>	The current status of the client on which the node resides*
<b>Client ID</b>	The ID of the client.*
<b>Host</b>	The name of the host.*
<b>Copyset</b>	The name of the Copyset hosted by the node.*
<b>PID</b>	The process ID of the node process.*
<b>Node Ready</b>	When checked, the node is operational.
<b>Node Started</b>	When checked, the node has been started and is up and running.*
<b>Messages Rcvd Rate</b>	The rate of messages received on the node.
<b>Messages Sent Rate</b>	The rate of messages sent on the node.
<b>CPU Used (%)</b>	The percentage of CPU memory used on the node.*
<b>Memory Used (kb)</b>	The amount of memory used by the node, in kilobytes.
<b>Get Count</b>	The total number of "get" operation performed on the node.*
<b>Get Rate</b>	The rate of "get" operations (per second) performed on the node.*
<b>Put Count</b>	The total number of "put" operations performed on the node.*
<b>Put Rate</b>	The rate of "put" operations (per second) performed on the node.*
<b>Remove Count</b>	The total number of "remove" operations performed on the node.*
<b>Remove Rate</b>	The rate of "remove" operations (per second) performed on the node.*
<b>Operations Completed</b>	The total number of operations completed on the node.*
<b>Operations Completed Rate</b>	The rate at which operations are completed on the node.*
<b>Operations Failed</b>	The number of operations failed on the node.*
<b>Operations Failed Rate</b>	The rate at which operations are failing on the node.*
<b>Operations Allowed</b>	The number of operations allowed on the node.*
<b>Operations Pending</b>	The number of operations pending on the node.*

<b>Operations Rejected</b>	The number of operations rejected on the node.*
<b>Operations Suspended</b>	The number of operations suspended on the node.*
<b>Iterator Create Count</b>	The number of iterator operations created on the node.*
<b>Iterator Create Rate</b>	The rate of iterator operations being created on the node.*
<b>Iterator Close Count</b>	The number of closed iterator operations on the node.*
<b>Iterator Close Rate</b>	The rate of iterator operations (per second) on the node.*
<b>Iterator Get Count</b>	The number of "get" iterator operations on the node.*
<b>Iterator Get Rate</b>	The rate of "get" iterator operations on the node.*
<b>Query Create Count</b>	The number of created queries on the node.*
<b>Query Create Rate</b>	The rate of created queries on the node.*
<b>Query Close Count</b>	The number of queries that were closed on the node.*
<b>Query Close Rate</b>	The rate of closed queries on the node.*
<b>Query Get Count</b>	The total number of "get" queries performed on the node.*
<b>Query Get Rate</b>	The rate of "get" queries (per second) performed on the node.*
<b>Transaction Begin Count</b>	The number of transactions started on the node.*
<b>Transaction Begin Rate</b>	The rate of transactions started on the node.*
<b>Transaction Commit Count</b>	The number of transactions committed on the node.*
<b>Transaction Commit Rate</b>	The rate of transactions committed on the node.*
<b>Transaction Rollback Count</b>	The number of transactions that have been rolled back on the node.*
<b>Transaction Rollback Rate</b>	The rate of transactions that have been rolled back on the node.*
<b>Statement Create Count</b>	The number of statements created on the node.*
<b>Statement Close Count</b>	The number of statements closed on the node.*
<b>Rows Expired</b>	The number of expired rows on the node.*

<b>Rows Received</b>	The number of received rows on the node.*
<b>Rows Scanned</b>	The number of scanned rows on the node.*
<b>Scans Completed</b>	The number of completed scans on the node.*
<b>Reindex Started</b>	The number of "reindex" operations started on the node.*
<b>Reindex Completed</b>	The number of "reindex" operations completed on the node.*
<b>Reindex Row Count</b>	The number of "reindex" rows on the node.*
<b>KVStore Write Count</b>	The number of "writes" to the KVStore on the node.*
<b>KVStore Write Time</b>	The amount of time taken to write to the KVStore on the node.*
<b>KVStore Value Count</b>	The values written to the key-value store.*
<b>Migration Bins Pending</b>	The number of "bins" awaiting migration.*
<b>Migration Bins Rcvd</b>	The number of "bins" received for migration.*
<b>Migration Bins Sent</b>	The number of "bins" sent to be migrated.*
<b>Deleted Keys</b>	The number of deleted keys on the node.*
<b>Events Forwarded</b>	The number of events that were forwarded on the node.*
<b>Listeners</b>	The number of listeners on the node.*
<b>Live Data Size</b>	The size of the live data.*
<b>TTL Tables</b>	The number of tables with row expiration defined.*
<b>Value Queries</b>	The number of value queries on the node.*
<b>Timestamp</b>	The date and time the row data was last updated.

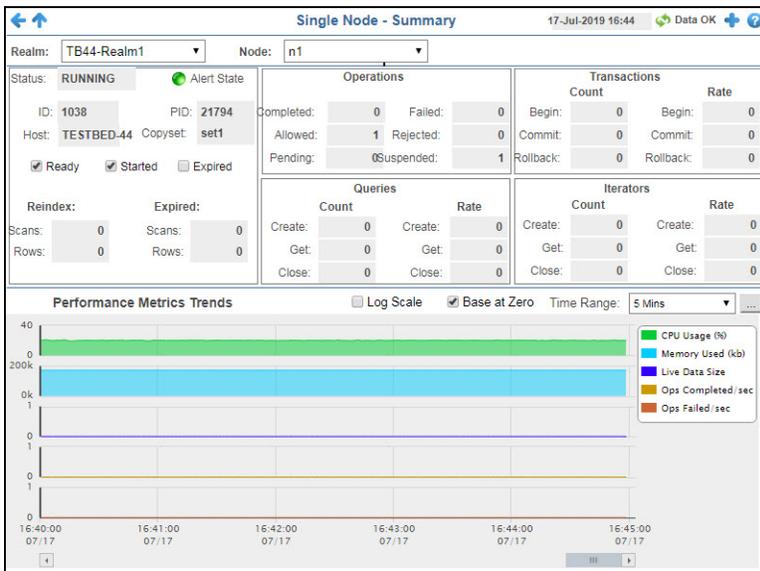
## Single Node Views

These displays provide detailed data for a single node as well as for node workers and node queues for a particular realm/node combination. Displays in this View are:

- **"Node Summary"**: This display allows you to view metrics and trend data for a particular node.

## Node Summary

This display provides a view of the current and historical metrics for a single node. The trend graph in the bottom half of the display traces the current and historical CPU usage, memory usage, live data size, rate of completed operations, and the rate of failed operations for the node.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu**, **Table** open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the selected domain. Refer to TIBCO ActiveSpaces documentation for more information regarding these fields.

### Filter By:

The display might include these filtering options:

- Realm** Select the realm (containing the node) for which you want to show data in the display.
- Node** Select the node for which you want to show data in the display.

### Fields and Data:

<b>Status</b>	The current status of the node.
<b>Alert State</b>	The current alert severity.  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
<b>ID</b>	The ID of the node.*
<b>Host</b>	The name of the host.*
<b>PID</b>	The process ID of the node.
<b>Copyset</b>	The name of the copyset.*
<b>Ready</b>	When checked, denotes that the node is ready and available.*
<b>Started</b>	When checked, denotes that the node is up and running.*
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > (Project Name) > <b>Solution Package Configuration</b> > <b>TIBCO Active Spaces</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Reindex</b>	<b>Scans</b> -- The number of "reindex" scans on the node. <b>Rows</b> -- The number of "reindex" rows on the node.
<b>Expired</b>	<b>Scans</b> -- The number of expired scans on the node. <b>Rows</b> -- The number of expired rows on the node.
<b>Operations</b>	
	<b>Completed</b> The number of completed operations on the node.
	<b>Allowed</b> The number of allowed operations on the node.
	<b>Pending</b> The number of pending operations on the node.
	<b>Failed</b> The number of failed operations on the node.
	<b>Rejected</b> The number of rejected operations on the node.
	<b>Suspended</b> The number of suspended operations on the node.
<b>Transactions</b>	
	<b>Count</b> <b>Begin</b> -- The number of transactions started on the node. <b>Commit</b> -- The number of transactions committed on the node. <b>Rollback</b> -- The number of transactions that have been rolled back on the node.
	<b>Rate</b> <b>Begin</b> -- The rate of transactions started on the node. <b>Commit</b> -- The rate of transactions committed on the node. <b>Rollback</b> -- The rate of transactions that have been rolled back on the node.
<b>Queries</b>	

- Count**      **Create** -- The number of created queries on the node.  
**Get** -- The number of "get" operations on the node.  
**Close** -- The number of closed queries on the node.
- Rate**      **Create** -- The rate of created queries on the node.  
**Get** -- The rate of "get" operations on the node.  
**Close** -- The rate of closed queries on the node.

### Iterators

- Count**      **Create** -- The number of iterator operations on the node.  
**Get** -- The number of "get" iterator operations on the node.  
**Close** -- The number of closed iterator operations on the node.
- Rate**      **Create** -- The rate of iterator operations on the node.  
**Get** -- The rate of "get" iterator operations on the node.  
**Close** -- The rate of closed iterator operations on the node.

### Performance Metric Trends

Traces the following:

- CPU Usage (%)** -- traces the CPU usage percentage for the node.  
**Memory Used (kb)**-- traces the amount of memory used, in kilobytes.  
**Live Data Size**-- traces the Live Data Size.  
**Ops Completed/s** -- traces the rate of completed operations.  
**Ops Failed/s** -- traces the rate of failed operations.

**Log Scale**      Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.

**Base at Zero**      Select to use zero (**0**) as the Y axis minimum for all graph traces.

**Time Range**      Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## All Proxy Views

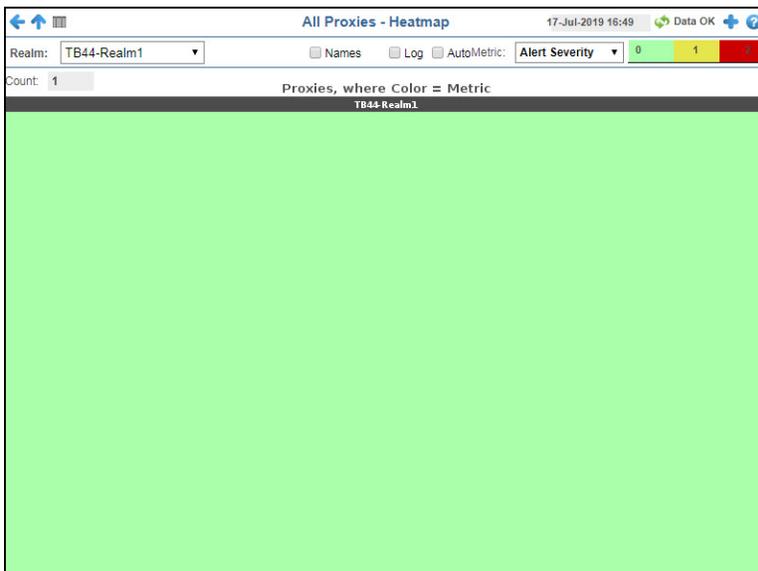
These displays provide detailed data for all proxies in a heatmap and tabular format. Displays in this View are:

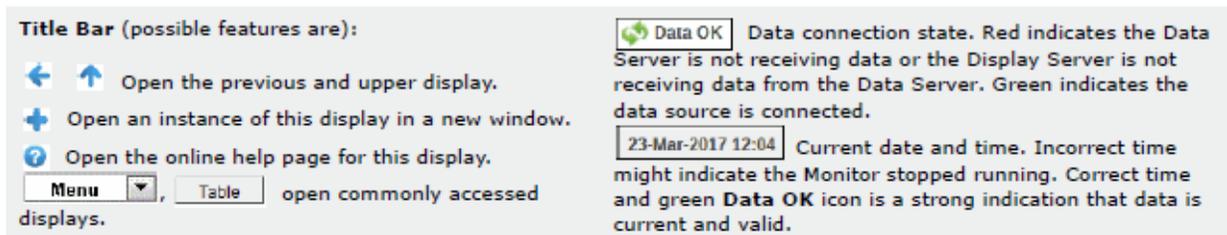
- [“All Proxies Heatmap”](#): A heatmap view of all proxies contained within a particular realm.
- [“All Proxies Table”](#): A tabular view of your proxies and their associated metrics within a particular realm.

### All Proxies Heatmap

This heatmap display provides an easy-to-view interface that allows you to quickly identify the current status of each proxy for each available metric. You can view the proxies in the heatmap based on the following metrics: current alert severity, alert count, CPU usage, memory used, rate of messages received, rate of messages sent, and transaction rollback rate. By default, this display shows the heatmap based on the **Alert Severity** metric.

You can use the **Names** check-box  to include or exclude labels in the heatmap, and you can mouse over a rectangle to see additional metrics for a proxy. Clicking one of the rectangles in the heatmap opens the [“Proxy Summary”](#) display, which allows you to see additional details for the selected proxy.



**Filter By:**

**Realm** Select the realm for which you want to see data.

**Fields and Data:**

**Count** Displays the total number of proxies displayed in the heatmap.

**Names** Select this check box to display the names of the proxies at the top of each rectangle in the heatmap.

**Log** Select this check box to use a logarithmic scale, rather than a linear scale, to map from the selected metric value for a cell to the color for the cell. **Log Scale** provides another way to distribute and differentiate values that you might not be able to see on a linear scale due to the dominant nature of large values in a linear scale.

**Auto** When checked, the values of the selected metric are auto-scaled to its highest defined value. When unchecked, the values of the selected metric display based on the threshold defined for the alert associated with the selected metric. Selecting Auto helps to visualize the range of the values currently present for the selected metric instead of the threshold of the alert that has been associated with the metric. All metrics that have not been associated in the heatmap defaults with alerts use a monochromatic color gradient bar (whites and greens). All metrics that have been associated in the heatmap defaults with alerts use a multi-chromatic color gradient bar (reds, yellows, white, and greens).

**Color Metric** Choose a metric to view in the display.

**Alert Severity**

The current alert severity. Values range from **0** - **2**, as indicated in the color gradient  bar, where **2** is the highest Alert Severity:

 Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.

 Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.

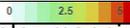
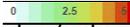
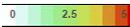
 Green indicates that no metrics have exceeded their alert thresholds.

**Alert Count**

The total number of alarm and warning alerts in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

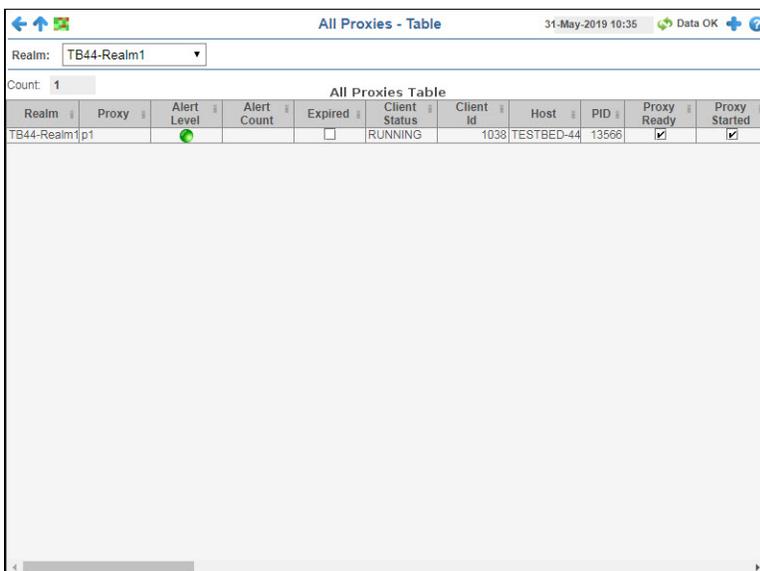
**CPU Usage**

The CPU usage rate for the proxy. The color gradient bar  , populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TdgProxyCpuUsageHigh**. The middle value in the gradient bar indicates the middle value of the range.

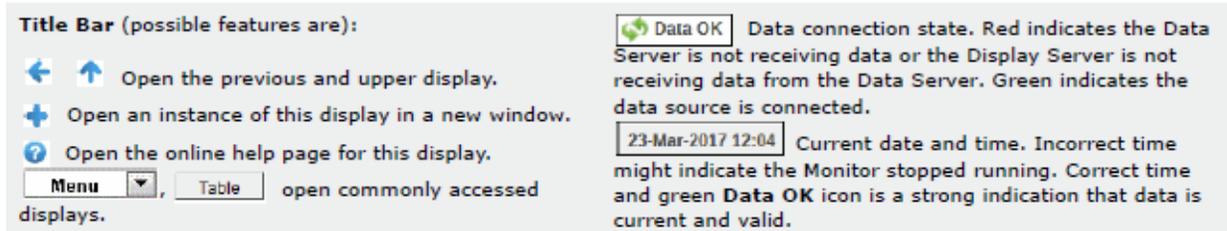
- Memory**      The memory usage for the proxy. The color gradient  , populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TdgProxyMemoryUseHigh**. The middle value in the gradient bar indicates the middle value of the range.
- Msgs Rcvd/sec**      The rate of messages received. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TdgProxyMsgsRcvdRateHigh**. The middle value in the gradient bar indicates the middle value of the range.
- Msgs Sent/sec**      The rate of messages sent. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TdgProxyMsgsSentRateLow**. The middle value in the gradient bar indicates the middle value of the range.
- Txn Rollback Rate**      The rate of rolled back transactions. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TdgProxyTxnRollbackRateHigh**. The middle value in the gradient bar indicates the middle value of the range.

## All Proxies Table

The table in this display provides a view of all proxies and their associated metric data in a selected realm. You can click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by clicking a row to view details for the selected proxy in the "Proxy Summary" display



Realm	Proxy	Alert Level	Alert Count	Expired	Client Status	Client Id	Host	PID	Proxy Ready	Proxy Started
TB44-Realm1	p1			<input type="checkbox"/>	RUNNING	1038	TESTBED-44	13566	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>




---

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the selected domain. Refer to TIBCO ActiveSpaces documentation for more information regarding these fields.

---

### Filter By:

<b>Realm</b>	Select the realm for which you want to view data.
<b>Count</b>	The total number of proxies found for the realm selected in the <b>Realm</b> dropdown, which are displayed in the <b>All Proxies Table</b> .

### All Proxies Table

<b>Realm</b>	The name of the realm.
<b>Proxy</b>	The name of the proxy.
<b>Alert Level</b>	The current alert severity. <ul style="list-style-type: none"> <li><span style="color: red;">●</span> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li><span style="color: yellow;">●</span> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li><span style="color: green;">●</span> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Alert Count</b>	The total number of alerts for the proxy.
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > (Project Name) > <b>Solution Package Configuration</b> > <b>TIBCO Active Spaces</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Client Status</b>	The status of the client.*
<b>Client ID</b>	The ID of the client.*
<b>Host</b>	The name of the host.*
<b>PID</b>	The process ID of the host.*
<b>Proxy Ready</b>	When checked, the proxy is up and running.*
<b>Proxy Started</b>	When checked, the proxy is up and running.*
<b>CPU Usage (%)</b>	The percentage of CPU used on the proxy.*

<b>Memory Used (kb)</b>	The amount of memory used, in kilobytes.*
<b>Get Count</b>	The total number of "get" operations performed on the proxy.*
<b>Get Rate</b>	The rate of "get" operations (per second) performed on the proxy.*
<b>Put Count</b>	The total number of "put" operations performed on the proxy.*
<b>Put Rate</b>	The rate of "put" operations (per second) performed on the proxy.*
<b>Remove Count</b>	The total number of "remove" operations performed on the proxy.*
<b>Remove Rate</b>	The rate of "remove" operations (per second) performed on the proxy.*
<b>Iterators</b>	The number of Iterators on the proxy.*
<b>Queries</b>	The number of queries on the proxy.*
<b>Statements</b>	The number of statements on the proxy.*
<b>Listeners</b>	The number of listeners on the proxy.*
<b>Iterator Create Count</b>	The number of iterator operations on the proxy.*
<b>Iterator Create Rate</b>	The rate of iterator operations being created on the proxy.*
<b>Iterator Close Count</b>	The number of iterator operations being closed on the proxy.*
<b>Iterator Close Rate</b>	The rate of iterator operations being closed on the proxy.*
<b>Iterator Get Count</b>	The number of "get" iterator operations on the proxy.*
<b>Iterator Get Rate</b>	The rate of "get" iterator operations on the proxy.*
<b>Query Create Count</b>	The number of created queries on the proxy.*
<b>Query Create Rate</b>	The rate of created queries on the proxy.*
<b>Query Close Count</b>	The number of closed queries on the proxy.*
<b>Query Close Rate</b>	The rate of closed queries on the proxy.*
<b>Query Get Count</b>	The number of "get" queries on the proxy.*
<b>Query Get Rate</b>	The rate of "get" queries on the proxy.*
<b>Transaction Begin Count</b>	The number of transactions started on the proxy.*
<b>Transaction Begin Rate</b>	The rate of transactions being started on the proxy.*
<b>Transaction Commit Count</b>	The number of commit transactions on the proxy.*

<b>Transaction Commit Rate</b>	The rate of transactions being committed on the proxy.*
<b>Transaction Rollback Count</b>	The number of transactions rolled back on the proxy.*
<b>Transaction Rollback Rate</b>	The rate of transactions rolled back on the proxy.*
<b>Client Connected</b>	The number of clients connected.*
<b>Client Disconnected</b>	The number of clients disconnected.*
<b>Timestamp</b>	The date and time the row data was last updated.

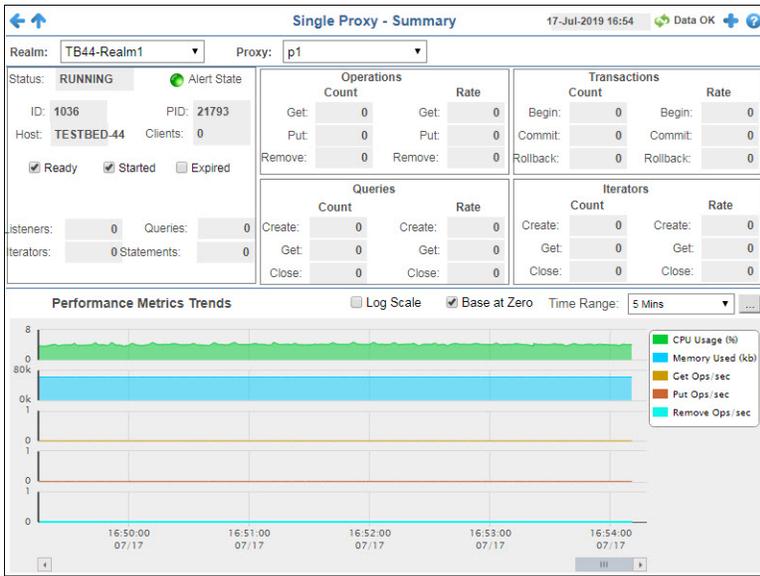
## Single Proxy Views

These displays provide detailed data for a single proxy in a summary format and list of all queues for a realm/proxy combination. Displays in this View are:

- **“Proxy Summary”**: This display allows you to view metrics and trend data for a particular proxy.

### Proxy Summary

This display provides a view of the current and historical metrics for a single proxy. The trend graph in the bottom half of the display traces the current and historical total number of CPU usage, memory usage, rate of get operations, rate of put operations, and the rate of remove operations.



**Title Bar (possible features are):**

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu Table open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the selected domain. Refer to TIBCO ActiveSpaces documentation for more information regarding these fields.

**Filter By:**

The display might include these filtering options:

**Realm** Select the realm (containing the proxy) for which you want to show data in the display.

**Proxy** Select the proxy for which you want to show data in the display.

**Fields and Data:**

**Status** The current status of the proxy.

**Alert State** The current alert severity.

- Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
- Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
- Green indicates that no metrics have exceeded their alert thresholds.

**ID** The ID of the proxy.\*

<b>Host</b>	The name of the host.*
<b>PID</b>	The process ID of the proxy.
<b>Clients</b>	The number of clients.*
<b>Ready</b>	When checked, denotes that the proxy is ready and available.*
<b>Started</b>	When checked, denotes that the proxy is up and running.*
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > (Project Name) > <b>Solution Package Configuration</b> > <b>TIBCO Active Spaces</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Listeners</b>	The total number of listeners on the proxy.
<b>Iterators</b>	The total number of iterators on the proxy.
<b>Queries</b>	The total number of queries on the proxy.
<b>Statements</b>	The total number of statements on the proxy.
<b>Operations</b>	
	<b>Count</b>
	<b>Get</b> -- The number of "get" operations on the proxy.
	<b>Put</b> -- The number of "put" operations on the proxy.
	<b>Remove</b> -- The number of "remove" operations on the proxy.
	<b>Rate</b>
	<b>Get</b> -- The rate of "get" operations on the proxy.
	<b>Put</b> -- The rate of "put" operations on the proxy.
	<b>Remove</b> -- The rate of "remove" operations on the proxy.
<b>Transactions</b>	
	<b>Count</b>
	<b>Begin</b> -- The number of transactions started on the proxy.
	<b>Commit</b> -- The number of transactions committed on the proxy.
	<b>Rollback</b> -- The number of transactions rolled back on the proxy.
	<b>Rate</b>
	<b>Begin</b> -- The rate of transactions being started on the proxy.
	<b>Commit</b> -- The rate of transactions being committed on the proxy.
	<b>Rollback</b> -- The rate of transactions being rolled back on the proxy.
<b>Queries</b>	
	<b>Count</b>
	<b>Create</b> -- The number of queries created on the proxy.
	<b>Get</b> -- The number of "get" queries on the proxy.
	<b>Close</b> -- The number of queries closed on the proxy.
	<b>Rate</b>
	<b>Create</b> -- The rate of queries being created on the proxy.
	<b>Get</b> -- The rate of "get" queries being created on the proxy.
	<b>Close</b> -- The rate of queries being closed on the proxy.
<b>Iterators</b>	
	<b>Count</b>
	<b>Create</b> -- The number of iterator operations created on the proxy.
	<b>Get</b> -- The number of "get" iterator operations on the proxy.
	<b>Close</b> -- The number of closed iterator operations on the proxy.

- Rate**
- Create** -- The rate of iterator operations being created on the proxy.
  - Get** -- The rate of "get" iterator operations being created on the proxy.
  - Close** -- The rate of iterator operations being closed on the proxy.

**Performance Metric Trends**

Traces the following:

- CPU Usage (%)** -- traces the percentage of CPU used for the node.
- Memory Used (kb)**-- traces the amount of memory used, in kilobytes.
- Get Ops/sec** -- traces the rate of "get" operations on the proxy.
- Put Ops/sec**-- traces the rate of "put" operations on the proxy.
- Remove Ops/sec** -- traces the rate of "remove" operations on the proxy.

**Log Scale** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.

**Base at Zero** Select to use zero (0) as the Y axis minimum for all graph traces.

**Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## All Keeper Views

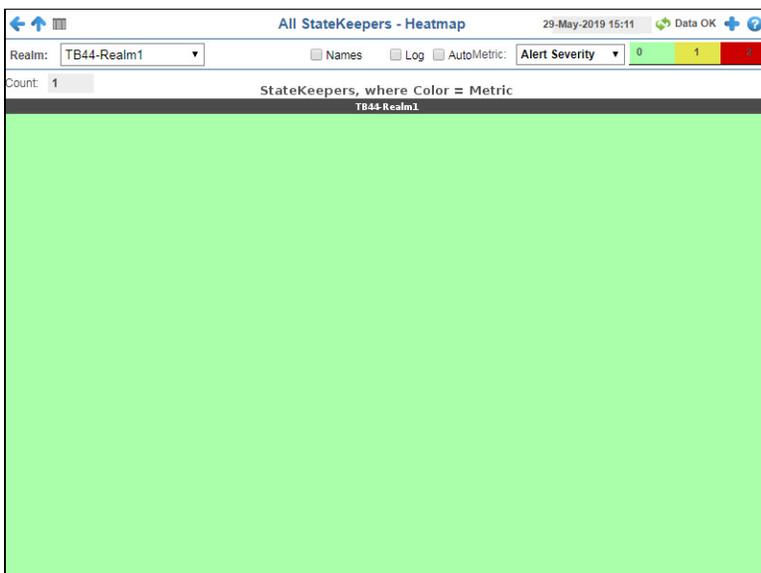
These displays provide detailed data for all keepers in a heatmap or tabular format. Displays in this View are:

- ["All Keepers Heatmap"](#): A heatmap view of all keepers contained within a particular realm.
- ["All Keepers Table"](#): A tabular view of all keepers and their associated metrics within a particular realm.

### All Keepers Heatmap

This heatmap display provides an easy-to-view interface that allows you to quickly identify the current status of each of your keepers for each available metric. You can view the keepers in the heatmap based on the following metrics: current alert severity, alert count, CPU usage, memory usage, and rate of messages received. By default, this display shows the heatmap based on the **Alert Severity** metric.

You can use the **Names** check-box  to include or exclude labels in the heatmap, and you can mouse over a rectangle to see additional metrics for a keeper. Clicking one of the rectangles in the heatmap opens the ["Keeper Summary"](#) display, which allows you to see additional details for the selected keeper.



#### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Filter By:**

**Realm** Select the realm for which you want to see data.

**Fields and Data:**

**Count** The total number of keepers found for the selected realm.

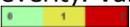
**Names** Select this check box to display the names of the keepers at the top of each rectangle in the heatmap.

**Log** Select this check box to use a logarithmic scale, rather than a linear scale, to map from the selected metric value for a cell to the color for the cell. **Log Scale** provides another way to distribute and differentiate values that you might not be able to see on a linear scale due to the dominant nature of large values in a linear scale.

**Auto** When checked, the values of the selected metric are auto-scaled to its highest defined value. When unchecked, the values of the selected metric display based on the threshold defined for the alert associated with the selected metric. Selecting Auto helps to visualize the range of the values currently present for the selected metric instead of the threshold of the alert that has been associated with the metric. All metrics that have not been associated in the heatmap defaults with alerts use a monochromatic color gradient bar (whites and greens). All metrics that have been associated in the heatmap defaults with alerts use a multi-chromatic color gradient bar (reds, yellows, white, and greens).

**Metric** Choose a metric to view in the display.

**Alert Severity**

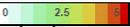
The current alert severity. Values range from **0** - **2**, as indicated in the color gradient  bar, where **2** is the highest Alert Severity:

-  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
-  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
-  Green indicates that no metrics have exceeded their alert thresholds.

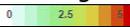
**Alert Count**

The total number of alarm and warning alerts in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

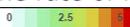
**CPU Usage**

The CPU usage rate for the keeper. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TdgKeeperCpuUsageHigh**. The middle value in the gradient bar indicates the middle value of the range.

**Memory**

The usage memory for the keeper. The color gradient bar  , populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TdgKeeperMemoryUseHigh**. The middle value in the gradient bar indicates the middle value of the range.

**Msgs Rcvd/sec**

The rate of messages received. The color gradient bar  , populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TdgKeeperMsgsRcvdRateHigh**. The middle value in the gradient bar indicates the middle value of the range.

## All Keepers Table

The table in this display provides a view of all keepers and their associated metric data for a specific realm. You can click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by clicking a row to view details for the selected keeper in the “Keeper Summary” display

Realm	Keeper	Alert Level	Alert Count	Expired	Client Status	Client Id	Host	Version	PID	Ready
TB44-Realm1	k1			<input type="checkbox"/>	RUNNING	1040	TESTBED-44		13565	<input checked="" type="checkbox"/>

### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** , **Table** open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the selected domain. Refer to TIBCO ActiveSpaces documentation for more information regarding these fields.

### Filter By:

- Realm** Select the realm for which you want to view data.
- Count** The total number of keepers found for the realm selected in the **Realm** dropdown, which are displayed in the **All StateKeepers Table**.

### All StateKeepers Table

<b>Realm</b>	The name of the realm.
<b>Keeper</b>	The name of the keeper.
<b>Alert Level</b>	The current alert severity.  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
<b>Alert Count</b>	The total number of alerts for the host.
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > (Project Name) > <b>Solution Package Configuration</b> > <b>TIBCO Active Spaces</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Client Status</b>	The current status of the client on which the keeper resides.*
<b>Client ID</b>	The ID of the client.*
<b>Host</b>	The name of the host.*
<b>PID</b>	The process ID of the StateKeeper process.*
<b>Ready</b>	When checked, the keeper is operational.*
<b>Started</b>	When checked, the keeper has been started and is up and running.*
<b>CPU Usage (%)</b>	The percentage of CPU memory used by the keeper.
<b>Memory Used (kb)</b>	The memory used by the keeper, in kilobytes.
<b>Message Rcvd Rate</b>	The rate of messages received.
<b>Message Sent Rate</b>	The rate of messages sent.
<b>Timestamp</b>	The date and time the row data was last updated.

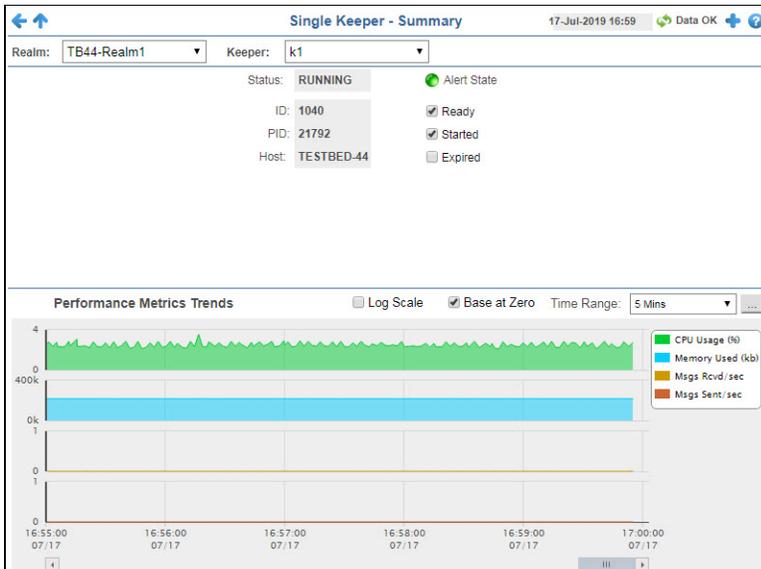
## Single Keeper Views

The display in this view provides detailed data for a single keeper in a summary format. Displays in this View are:

- **"Keeper Summary"**: This display allows you to view metrics and trend data for a particular keeper.

## Keeper Summary

This display provides a view of the current and historical metrics for a single keeper. The trend graph in the bottom half of the display traces the current and historical CPU usage, memory usage, rate of received messages, and the rate of sent messages for the keeper.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu, Table open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the selected domain. Refer to TIBCO ActiveSpaces documentation for more information regarding these fields.

### Filter By:

The display might include these filtering options:

- Realm** Select the realm (containing the keeper) for which you want to show data in the display.
- Keeper** Select the keeper for which you want to show data in the display.

### Fields and Data:

- Status** The current status of the keeper.\*
- ID** The ID of the keeper.\*

<b>PID</b>	The process ID of the StateKeeper.*
<b>Host</b>	The name of the host.*
<b>Alert State</b>	The current alert severity. <span style="color: red;">●</span> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold. <span style="color: yellow;">●</span> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold. <span style="color: green;">●</span> Green indicates that no metrics have exceeded their alert thresholds.
<b>Ready</b>	When checked, the keeper is operational.*
<b>Started</b>	When checked, the keeper has been started and is up and running.*
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > (Project Name) > <b>Solution Package Configuration</b> > <b>TIBCO Active Spaces</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Trends</b>	Traces the following: <b>CPU Usage (%)</b> -- traces the CPU usage percentage. <b>Memory Used (kb)</b> -- traces the memory usage, in kilobytes. <b>Msgs Rcvd/sec</b> -- traces the rate of messages received, per second. <b>Msgs Sent/sec</b> -- traces the rate of messages sent, per second.
<b>Log Scale</b>	Select to enable a logarithmic scale. Use <b>Log Scale</b> to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. <b>Log Scale</b> makes data on both scales visible by applying logarithmic values rather than actual values to the data.
<b>Base at Zero</b>	Select to use zero ( <b>0</b> ) as the Y axis minimum for all graph traces.
<b>Time Range</b>	Select a time range from the drop down menu varying from <b>2 Minutes</b> to <b>Last 7 Days</b> , or display <b>All Data</b> . To specify a time range, click Calendar <input type="button" value="..."/> .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

---

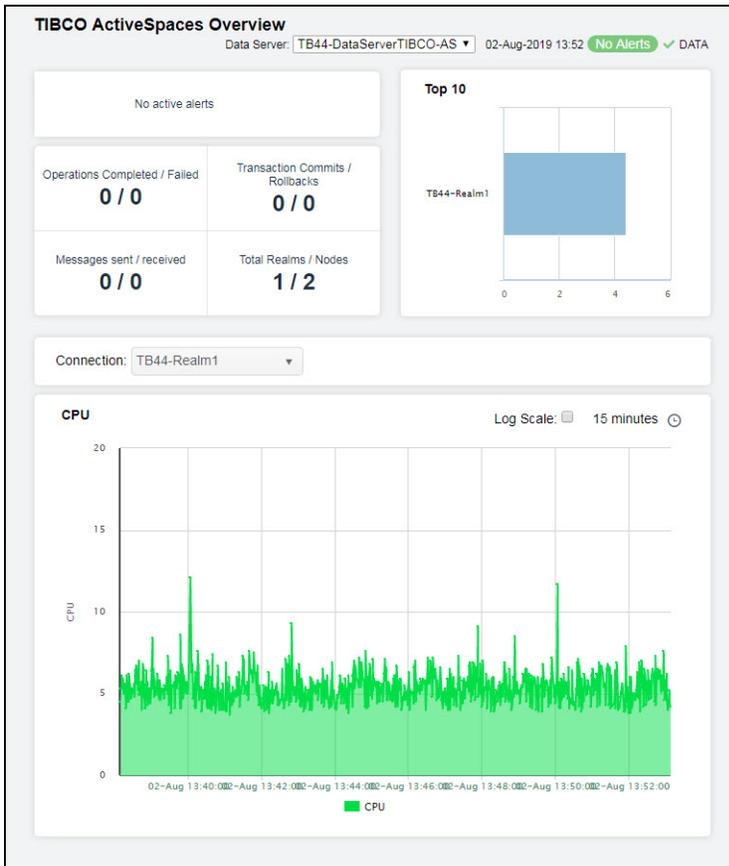
## TIBCO ActiveSpaces - HTML

The **TIBCO ActiveSpaces Overview** is the top-level display for the TIBCO ActiveSpaces Monitor, which provides a good starting point for immediately getting the status of all your operations, transactions, messages, and realms on your Data Server. You can select the RTView DataServer for which you want to see data and easily view the current data for that DataServer including:

- The total number of active alerts for the selected DataServer, including the total number of critical and warning alerts.
- The current number of operations completed and failed.
- The number of transactions committed and rolled back.
- The number of messages sent and received.
- The total number of realms and nodes.
- A visual list of the top 10 realms containing the total operations/messages/transactions/ realms on your connected DataServer.

You can hover over each region in the upper half of the Overview to see more detail. You can also drill down to see even more detail by clicking on each respective region in the Overview. For example, clicking on the alerts in the **CRITICAL** and **WARNING** alerts region opens the **Alerts Table by Components** display.

The bottom half of the display provide a **CPU** trend graph representing CPU usage percentage for a selected connection. You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.



The following TIBCO ActiveSpaces Views can be found under **Components** tab > **Middleware** > **TIBCO ActiveSpaces**:

- **"Realms Views - HTML"**: The displays in this View provide detailed data for all realms in a heatmap and tabular format, or for a particular realm in tabular and trend graph format.
- **"Nodes Views - HTML"**: The displays in this View provide detailed data for all nodes (in a specific realm) in a heatmap or tabular format.
- **"Proxies Views - HTML"**: The displays in this View provide detailed data for all proxies in a heatmap and tabular format.
- **"Keepers Views - HTML"**: The displays in this View provide detailed data for all keepers in a heatmap or tabular format.

## Realms Views - HTML

These displays provide detailed data for all realms in a heatmap and tabular format. Clicking Realms in the left/navigation menu opens the ["TIBCO ActiveSpaces Realms Table - HTML"](#), which provides a tabular view of your realms and their associated metrics. Displays in this View are:

- **All Realms Heatmap:** Opens the ["TIBCO ActiveSpaces Realms Heatmap - HTML"](#) display, which provides a heatmap view of all realms.
- **Single Realm Summary:** Opens the ["TIBCO ActiveSpaces Realm Summary - HTML"](#) display, which provides a view of the current and historical metrics for a single realm.
- **Realm Servers:** Opens the ["TIBCO ActiveSpaces Realm Servers - HTML"](#) display, which provides a view of your group server metrics, the group metrics, the persist server metrics, and the persist metrics for a selected realm.

## TIBCO ActiveSpaces Realms Table - HTML

The table in this display provides a view of all of your realms and their associated metric data including alert level, alert count, and the current value of each gathered metric. You can click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by double-clicking a row to view details for the selected realm in the ["TIBCO ActiveSpaces Realm Summary - HTML"](#) display.

TIBCO ActiveSpaces Realms Table 01-Jul-2019 13:33 No Alerts ✓ DATA

Count: 1

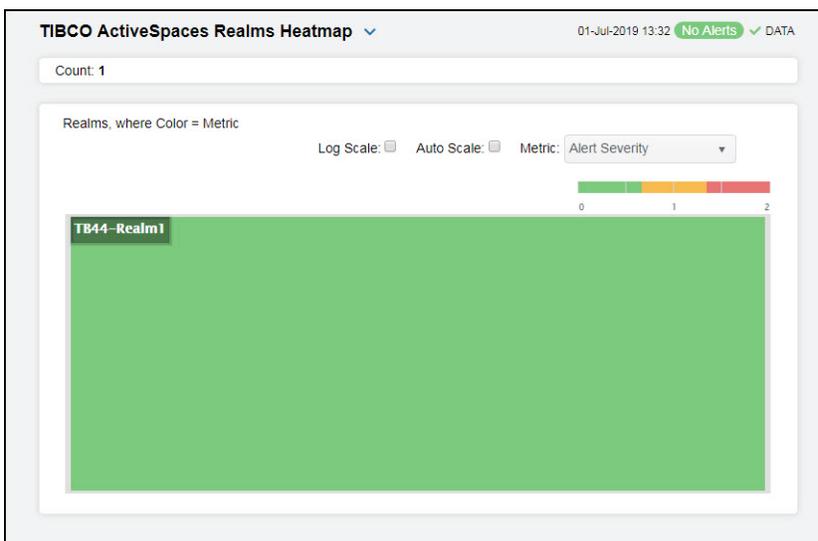
All Realms Table

Realm	Alert Level	Alert Count	Expired	Grid Name	Server Cpu	Server Memory	Server ID
TB44-Realm1	✓			_default	6.600	529104	e58808d6-9a36-28 r

## TIBCO ActiveSpaces Realms Heatmap - HTML

Clicking **All Realms Heatmap** in the left/navigation menu opens the **TIBCO ActiveSpaces Realms Heatmap**, which provides an easy-to-view interface that allows you to quickly identify the current status of each of your realms for each available metric. You can view the realms in the heatmap based on the following metrics: current alert severity, alert count, CPU usage, memory usage, operations completed, and operations failed. By default, this display shows the heatmap based on the **Alert Severity** metric.

The heatmap is organized so that each rectangle represents a space. The rectangle color indicates the most critical alert state. Click on a node to drill-down to the "[TIBCO ActiveSpaces Realm Summary - HTML](#)" display and view metrics for a particular realm. Toggle between the commonly accessed displays by clicking the drop down list on the display title. Mouse-over rectangles to view more details about realm performance and status.



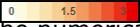
### Available Metrics

Select the metric driving the heatmap display. The default is **Alert Severity**. Each **Metric** has a color gradient bar that maps values to colors. The heatmap is organized by realms, where each rectangle represents a realm. Mouse-over any rectangle to display the current values of the metrics for the realm. Click on a rectangle to drill-down to the associated "[TIBCO ActiveSpaces Realm Summary - HTML](#)" display for a detailed view of metrics for that particular realm.

**Alert Severity** The current alert severity. Values range from **0** - **2**, as indicated in the color gradient  bar, where **2** is the highest Alert Severity:

-  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
-  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
-  Green indicates that no metrics have exceeded their alert thresholds.

**Alert Count** The total number of alarm and warning alerts in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

- CPU Usage** The milliseconds of CPU time accumulated by the process after each update interval. The color gradient bar , populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TdgRealmServerCpuUsageHigh**. The middle value in the gradient bar indicates the middle value of the range.
- Memory** The amount of memory used in the realm. The color gradient bar , populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TdgRealmServerMemoryUseHigh**. The middle value in the gradient bar indicates the middle value of the range.
- Ops Completed** The number of operations completed in the realm. The color gradient bar , populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of operations in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
- Ops Failed** The number of failed operations in the realm. The color gradient bar , populated by the current heatmap, shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of operations that have failed in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

## TIBCO ActiveSpaces Realm Summary - HTML

Clicking **Single Realm Summary** in the left/navigation menu opens the **TIBCO ActiveSpaces Realm Summary** display, which provides a view of the current and historical metrics for a single realm. Hover over the boxes at the top of the display to view additional information. In the trend graph region, you can view the rate of completed operations, the rate of failed operations, the rate of transactions that are committed, and the rate of transactions that are rolled back over a selected time range.



## TIBCO ActiveSpaces Realm Servers - HTML

Clicking **Realm Servers** in the left/navigation menu opens the **TIBCO ActiveSpaces Realm Servers** display, which provides a view of your group server metrics, the group metrics, the persist server metrics, and the persist metrics for a selected realm. You can click a column header in each of the tables to sort column data in numerical or alphabetical order.

**TIBCO ActiveSpaces Realm Servers** 01-Jul-2019 13:45 No Alerts DATA

Realm: TB44-Realm1

---

**Group Server**

Connection	Group
TB44-Realm1	1036

---

**Group Metrics**

Connection	Group	Metric
TB44-Realm1	1036	FORMAT_UNAVAILABLE
TB44-Realm1	1036	DATA_LOST
TB44-Realm1	1036	BYTES_RECEIVED
TB44-Realm1	1036	BYTES_SENT
TB44-Realm1	1036	QUEUE_DISCARDS
TB44-Realm1	1036	QUEUE_BACKLOG
TB44-Realm1	1036	PROCESS_VM_KB
TB44-Realm1	1036	PROCESS_PEAK_RSS_KB

---

**Persist Server**

Connection	Persist
TB44-Realm1	2965

---

**Persist Metrics**

Connection	Persist	Metric
TB44-Realm1	tibdg_AC439694-761E-470F-9123-D8ADF5	MESSAGE_SIZE
TB44-Realm1	tibdg_AC439694-761E-470F-9123-D8ADF5	MESSAGE_COUNT
TB44-Realm1	tibdg_AC439694-761E-470F-9123-D8ADF5	_StoreDispatcher Event Queue.QUEUE_BA
TB44-Realm1	tibdg_AC439694-761E-470F-9123-D8ADF5	_StoreDispatcher Event Queue.QUEUE_DI
TB44-Realm1	tibdg_AC439694-761E-470F-9123-D8ADF5	DYNAMIC_FORMATS
TB44-Realm1	tibdg_AC439694-761E-470F-9123-D8ADF5	USER_CPU_TIME

⏪ ⏩ Page 1 of 2 ⏪ ⏩ 1 - 40 of 57 items

## Nodes Views - HTML

These displays provide detailed data for all nodes (in a specific realm) in a heatmap or tabular format. Clicking **Nodes** in the left/navigation menu opens the ["TIBCO ActiveSpaces Nodes Table - HTML"](#), which provides a tabular view of all nodes (contained within a particular realm) and their associated metrics. Displays in this View are:

- **All Nodes Heatmap:** Opens the ["TIBCO ActiveSpaces Nodes Heatmap - HTML"](#) display, which is a heatmap view of all nodes contained within a particular realm.
- **Single Node Summary:** Opens the ["TIBCO ActiveSpaces Node Summary - HTML"](#) display, which allows you to view metrics and trend data for a particular node.

## TIBCO ActiveSpaces Nodes Table - HTML

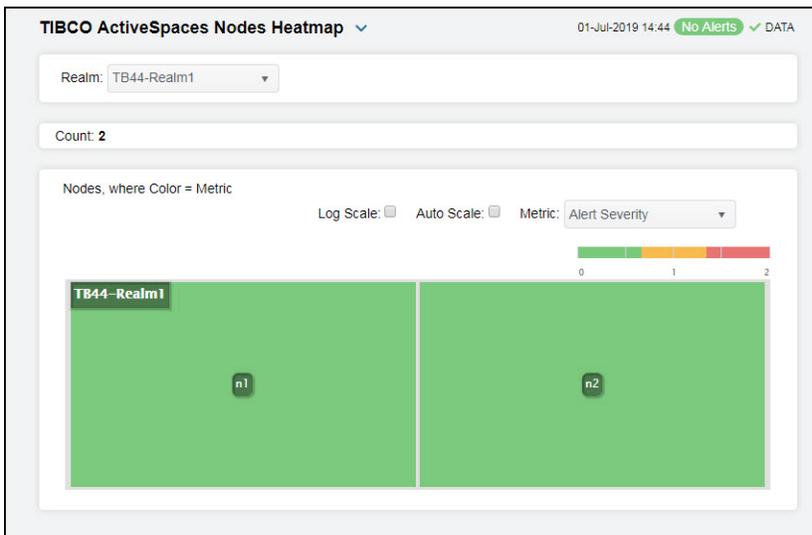
The table in this display provides a view of all nodes and their associated metric data in a specific realm. You can click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by double-clicking a row to view details for the selected node in the ["TIBCO ActiveSpaces Node Summary - HTML"](#) display.

Realm	Node	Alert Level	Alert Count	Expired	CPU Used	CPU Used/s	Memory Used	Me Us
TB44-Realm1	n1	✓			348817593	200.000		
TB44-Realm1	n2	✓			348878355	200.493		

## TIBCO ActiveSpaces Nodes Heatmap - HTML

Clicking **All Nodes Heatmap** in the left/navigation menu opens the **TIBCO ActiveSpaces Nodes Heatmap** display, which provides an easy-to-view interface that allows you to quickly identify the current status of each of your nodes for each available metric. You can view the nodes in the heatmap based on the following metrics: current alert severity, alert count, CPU usage, memory usage, rate of failed operations, and rate of completed operations. By default, this display shows the heatmap based on the **Alert Severity** metric.

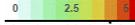
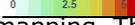
You can mouse over a rectangle to see additional metrics for a node. Clicking one of the rectangles in the heatmap opens the ["TIBCO ActiveSpaces Node Summary - HTML"](#) display, which allows you to see additional details for the selected node.



### Available Metrics

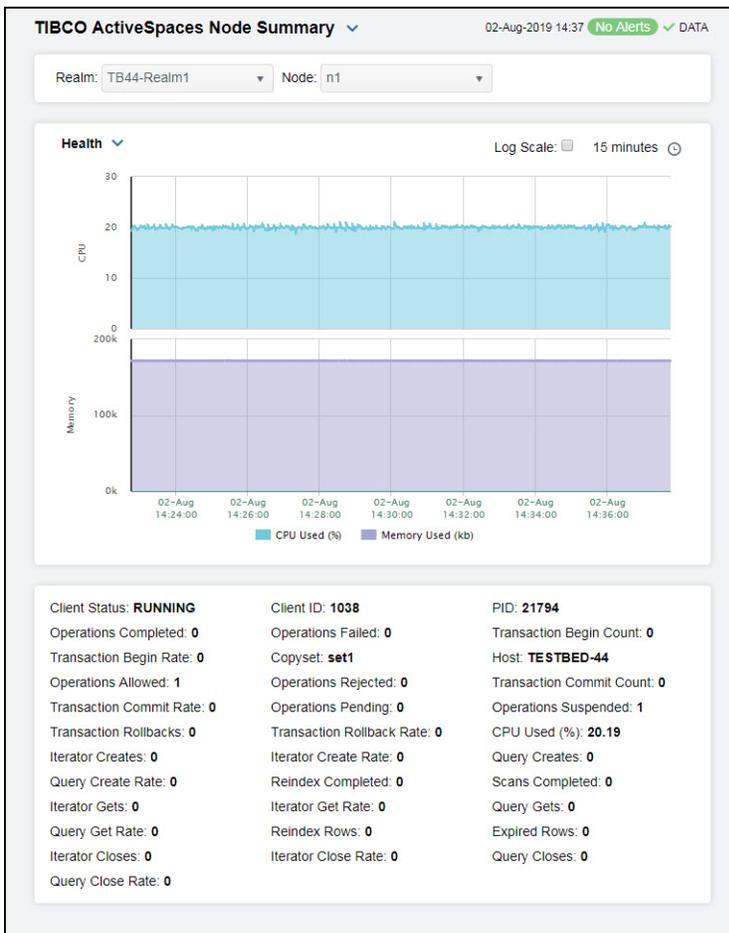
Select the metric driving the heatmap display. The default is **Alert Severity**. Each **Metric** has a color gradient bar that maps values to colors. The heatmap is organized by nodes, where each rectangle represents a node. Mouse-over any rectangle to display the current values of the metrics for the node. Click on a rectangle to drill-down to the associated ["TIBCO ActiveSpaces Node Summary - HTML"](#) display for a detailed view of metrics for that particular node.

- Alert Severity** The current alert severity. Values range from **0 - 2**, as indicated in the color gradient bar, where **2** is the highest Alert Severity:
- Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
  - Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
  - Green indicates that no metrics have exceeded their alert thresholds.
- Alert Count** The total number of alarm and warning alerts in a given item (index) associated with the rectangle. The color gradient bar shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
- CPU Usage** The milliseconds of CPU time accumulated by the process after the last update interval. The color gradient bar shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TdgNodeCpuUsageHigh**. The middle value in the gradient bar indicates the middle value of the range.

<b>Memory</b>	The memory usage for the node. The color gradient bar  , populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of <b>TdgNodeMemoryUseHigh</b> . The middle value in the gradient bar indicates the middle value of the range.
<b>Failed Op Rate</b>	The rate of failed operations. The color gradient bar  , populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of <b>TdgNodeOpsFailedRateHigh</b> . The middle value in the gradient bar indicates the middle value of the range.
<b>Completed Op Rate</b>	The rate of completed operations. The color gradient bar  , populated by the current heatmap, shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of <b>TdgNodeOpsCompletedRateLow</b> . The middle value in the gradient bar indicates the middle value of the range.

## TIBCO ActiveSpaces Node Summary - HTML

Clicking **Single Node Summary** in the left/navigation menu opens the **TIBCO ActiveSpaces Node Summary** display, which provides a view of the current and historical metrics for a single node. The trend graph in the bottom half of the display has three options: **Health**, **Live Data**, and **Operations**. **Health** traces the current and historical CPU usage and memory usage over a selected time range. **Live Data** traces the live data size over a selected time range. **Operations** traces the rate of completed operations and the rate of failed operations for the node over a selected time range.



## Proxies Views - HTML

These displays provide detailed data for all proxies (in a specific realm) in a heatmap or tabular format. Clicking **Proxies** in the left/navigation menu opens the ["TIBCO ActiveSpaces Proxies Table - HTML"](#) display, which provides a tabular view of your proxies and their associated metrics within a particular realm. Displays in this View are:

- **All Proxies Heatmap:** Opens the ["TIBCO ActiveSpaces Proxies Heatmap - HTML"](#) display, which provides a heatmap view of all proxies contained within a particular realm.
- **Single Proxy Summary:** Opens the ["TIBCO ActiveSpaces Proxy Summary - HTML"](#) display, which allows you to view metrics and trend data for a particular proxy.

## TIBCO ActiveSpaces Proxies Table - HTML

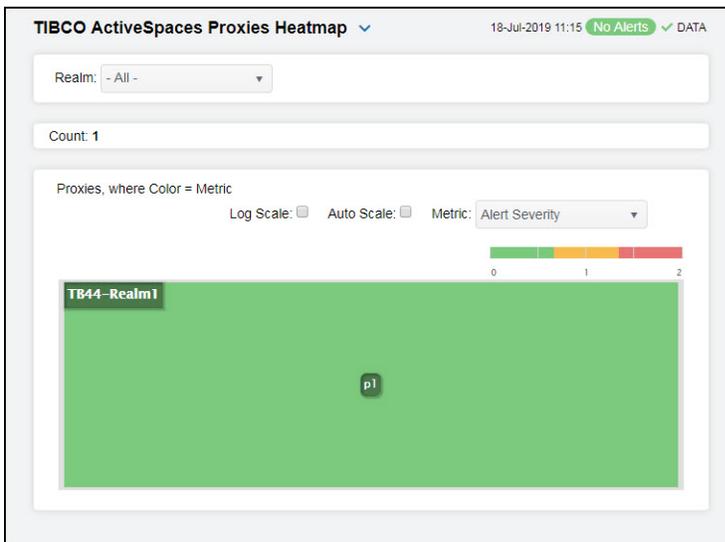
The table in this display provides a view of all proxies and their associated metric data in a selected realm. You can click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by double-clicking a row to view details for the selected proxy in the ["TIBCO ActiveSpaces Proxy Summary - HTML"](#) display

Realm	Proxy	Alert Level	Alert Count	Expired	CPU Used	CPU Used/s	Pr V
TB44-Realm1	p1	✓					

## TIBCO ActiveSpaces Proxies Heatmap - HTML

Clicking **All Proxies Heatmap** in the left/navigation menu opens the **TIBCO ActiveSpaces Proxies Heatmap**, which provides an easy-to-view interface that allows you to quickly identify the current status of each proxy for each available metric. You can view the proxies in the heatmap based on the following metrics: current alert severity, alert count, CPU usage, memory used, iterator count, listener count, query count, and statement count. By default, this display shows the heatmap based on the **Alert Severity** metric.

You can mouse over a rectangle to see additional metrics for a proxy. Clicking one of the rectangles in the heatmap opens the ["TIBCO ActiveSpaces Proxy Summary - HTML"](#) display, which allows you to see additional details for the selected proxy.



## Available Metrics

Select the metric driving the heatmap display. The default is **Alert Severity**. Each **Metric** has a color gradient bar that maps values to colors. The heatmap is organized by proxies, where each rectangle represents a proxy. Mouse-over any rectangle to display the current values of the metrics for the proxy. Click on a rectangle to drill-down to the associated ["TIBCO ActiveSpaces Proxy Summary - HTML"](#) display for a detailed view of metrics for that particular proxy.

**Alert Severity** The current alert severity. Values range from **0 - 2**, as indicated in the color gradient bar, where **2** is the highest Alert Severity:

- Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
- Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
- Green indicates that no metrics have exceeded their alert thresholds.

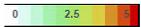
**Alert Count** The total number of alarm and warning alerts in a given item (index) associated with the rectangle. The color gradient bar shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

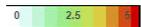
**CPU Usage** The CPU usage rate for the proxy. The color gradient bar , populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TdgProxyCpuUsageHigh**. The middle value in the gradient bar indicates the middle value of the range.

**Memory** The memory usage for the proxy. The color gradient bar , populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TdgProxyMemoryUseHigh**. The middle value in the gradient bar indicates the middle value of the range.

**Iterator Count** The number of iterators on the proxy. The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of iterators in the proxy. The middle value in the gradient bar indicates the middle value of the range.

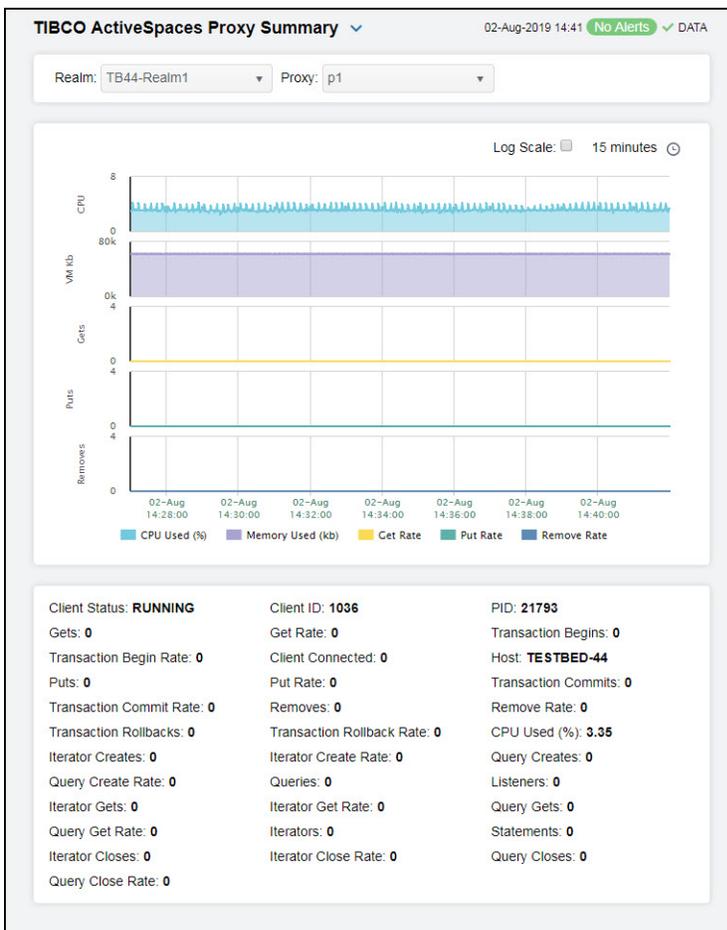
**Listener Count** The number of listeners on the proxy. The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of listeners in the proxy. The middle value in the gradient bar indicates the middle value of the range.

**Query Count** The number of queries on the proxy. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of queries in the proxy. The middle value in the gradient bar indicates the middle value of the range.

**Statement Count** The number of statements on the proxy. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of statements in the proxy. The middle value in the gradient bar indicates the middle value of the range.

### TIBCO ActiveSpaces Proxy Summary - HTML

Clicking **Single Proxy Summary** in the left/navigation menu opens the **TIBCO ActiveSpaces Proxy Summary**, which provides a view of the current and historical metrics for a single proxy. The trend graph in the display traces the current and historical rate of CPU usage, process virtual memory usage, rate of get operations, rate of put operations, and the rate of remove operations.



## Keepers Views - HTML

These displays provide detailed data for all keepers in a heatmap or tabular format, as well as metrics and trend data for a particular keeper. Clicking **Keepers** in the left/navigation menu opens the “[TIBCO ActiveSpaces StateKeepers Table - HTML](#)” display, which provides a tabular view of all keepers and their associated metrics within a particular realm. Displays in this View are:

- **All Keepers Heatmap:** Opens the “[TIBCO ActiveSpaces StateKeepers Heatmap - HTML](#)” display, which is a heatmap view of all keepers contained within a particular realm.
- **Single Keeper Summary:** Opens the “[TIBCO ActiveSpaces Keeper Summary - HTML](#)” display, which allows you to view metrics and trend data for a particular keeper.

## TIBCO ActiveSpaces StateKeepers Table - HTML

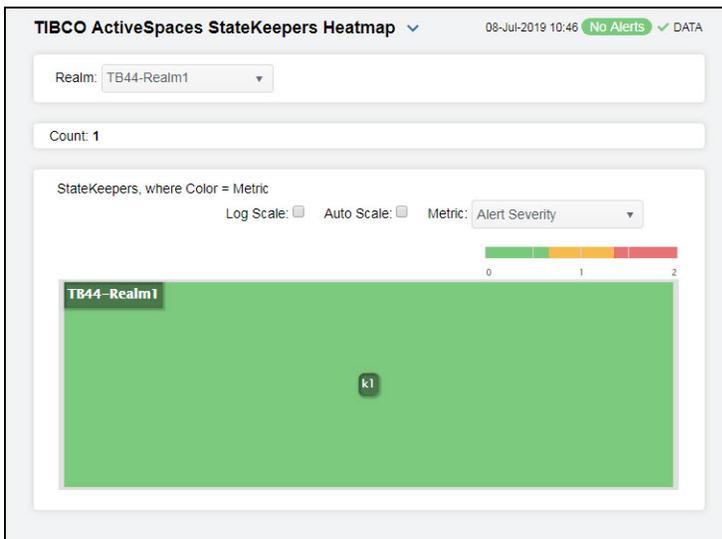
The table in this display provides a view of all keepers and their associated metric data for a specific realm. You can click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by double-clicking a row to view details for the selected keeper in the “[TIBCO ActiveSpaces Keeper Summary - HTML](#)” display

Realm	Keeper	Alert Level	Alert Count	Expired	CPU Used	CPU Used/s	Proc. Memor.
TB44-Realm1	k1	✓			10672510	45.854	

## TIBCO ActiveSpaces StateKeepers Heatmap - HTML

Clicking **All Keepers Heatmap** in the left/navigation menu opens the **TIBCO ActiveSpaces StateKeeper Heatmap**, which provides an easy-to-view interface that allows you to quickly identify the current status of each of your keepers for each available metric. You can view the keepers in the heatmap based on the following metrics: current alert severity, alert count, CPU usage, memory usage, rate of messages received, and rate of messages sent. By default, this display shows the heatmap based on the **Alert Severity** metric.

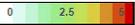
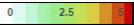
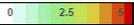
You can mouse over a rectangle to see additional metrics for a keeper. Clicking one of the rectangles in the heatmap opens the ["TIBCO ActiveSpaces Keeper Summary - HTML"](#) display, which allows you to see additional details for the selected keeper.



### Available Metrics

Select the metric driving the heatmap display. The default is **Alert Severity**. Each **Metric** has a color gradient bar that maps values to colors. The heatmap is organized by keepers, where each rectangle represents a keeper. Mouse-over any rectangle to display the current values of the metrics for the keeper. Click on a rectangle to drill-down to the associated ["TIBCO ActiveSpaces Keeper Summary - HTML"](#) display for a detailed view of metrics for that particular keeper.

- Alert Severity** The current alert severity. Values range from **0** - **2**, as indicated in the color gradient bar, where **2** is the highest Alert Severity:
- Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
  - Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
  - Green indicates that no metrics have exceeded their alert thresholds.
- Alert Count** The total number of alarm and warning alerts in a given item (index) associated with the rectangle. The color gradient bar shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
- CPU Usage** The CPU usage rate for the keeper. The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TdgKeeperCpuUsageHigh**. The middle value in the gradient bar indicates the middle value of the range.

- Memory** The usage memory for the keeper. The color gradient bar  , populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TdgKeeperMemoryUseHigh**. The middle value in the gradient bar indicates the middle value of the range.
- Msgs Rcvd/sec** The rate of messages received. The color gradient bar  , populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TdgKeeperMsgsRcvdRateHigh**. The middle value in the gradient bar indicates the middle value of the range.
- Msgs Sent/sec** The rate of messages received. The color gradient bar  , populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TdgKeeperMsgsSentRateLow**. The middle value in the gradient bar indicates the middle value of the range.

## TIBCO ActiveSpaces Keeper Summary - HTML

Clicking **Single Keeper Summary** in the left/navigation menu opens the **TIBCO ActiveSpaces Keeper Summary**, which provides a view of the current and historical metrics for a single keeper. The trend graph in the display traces the current and historical CPU usage rate, process memory usage, rate of received messages, and the rate of sent messages for the keeper.



---

## TIBCO ActiveSpaces (2.x)

The following TIBCO ActiveSpaces Views can be found under **Components** tab > **Middleware** > **TIBCO ActiveSpaces 2**:

- **"Spaces View"**: The displays in this View allow you to view the current and historical metrics for all metaspaces and spaces in a heatmap, tabular, or summary format.
- **"Members View"**: The displays in this View allow you to view the current and historical metrics for all members in a particular metaspaces, view data for members within a particular space, and view data for all spaces for a particular member.

### Spaces View

These displays provide detailed data for all metaspaces and spaces in a heatmap, tabular, or summary format. Displays in this View are:

- **"All Metaspaces Table"**: A tabular view of your metaspaces and their associated metrics.
- **"Metaspaces Summary"**: This display allows you to view metrics and trend data for a particular metaspaces.
- **"All Spaces Table"**: A tabular view of all spaces contained within a particular metaspaces.
- **"All Spaces Heatmap"**: A heatmap view of all spaces contained within a particular metaspaces.
- **"Space Summary"**: This display allows you to view metrics and trend data for a particular space.
- **"All Queries Table"**: This display allows you to view queries by domain, metaspaces, and space and view the performance metrics for the queries.
- **"Query Summary"**: This display allows you to view performance metrics for a particular query, as well as to view any related queries.

### All Metaspaces Table

The table in this display provides a view of all of your metaspaces and their associated metric data including domain, metaspaces, alert level, alert count, and the current value of each gathered metric. You can click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by clicking a row to view details for the selected metaspaces in the **"Metaspaces Summary"** display

ActiveSpaces All Metaspaces - Table 21-Apr-2017 11:48 Data OK

Domain: Production

Metaspace Count: 3

Domain	Metaspace	Alert Level	Alert Count	Spaces	Members	AS Version	Entries	Replicas
Production	retailV212		0	3	4	2.1.2.005	450	
Production	retailV213		0	3	5	2.1.3.005	469	
Production	retailV214		0	3	4	2.1.4.011	441	

**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu Table open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the selected domain. Refer to TIBCO ActiveSpaces documentation for more information regarding these fields.

**Filter By:**

- Domain** Select the domain for which you want to view data.
- Metaspace Count** The total number of metaspaces found for the domain selected in the **Domain** dropdown, which are displayed in the **Metaspaces** table.

**Metaspaces Table**

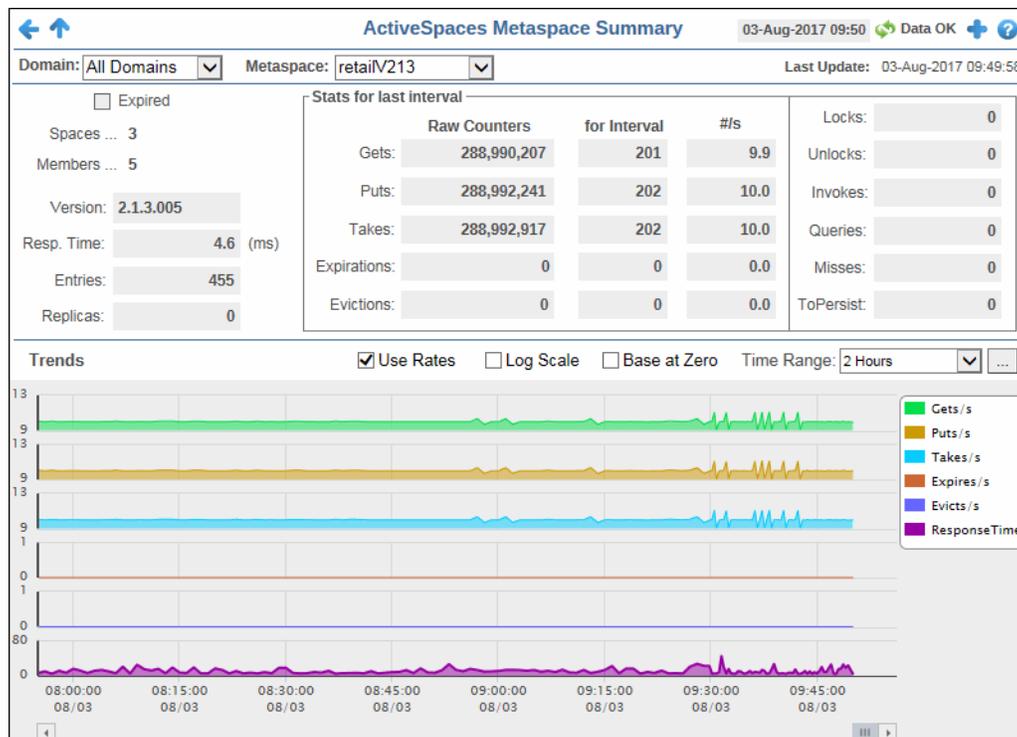
- Domain** The name of the domain.
- Metaspace** The name of the metaspace.

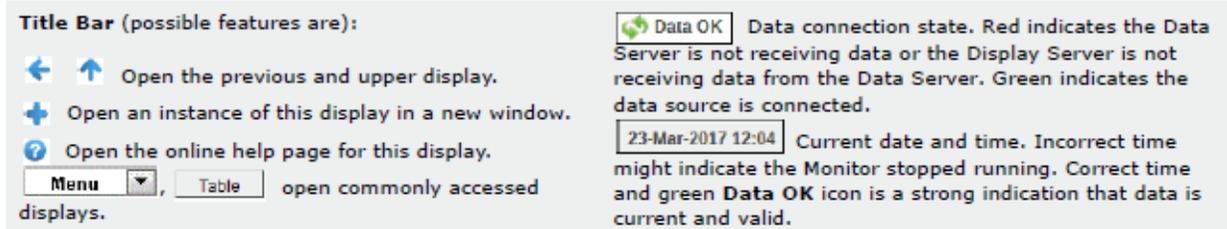
<b>Alert Level</b>	The current alert severity.  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
<b>Alert Count</b>	The total number of alerts for the host.
<b>Spaces</b>	The number of user spaces defined in the metaspace.*
<b>Members</b>	The number of members (clients and servers) associated with the metaspace.*
<b>AS Version</b>	The metaspace's current version of TIBCO ActiveSpaces.*
<b>Entries</b>	The total number of entries stored in the metaspace.*
<b>Replicas</b>	The total number of replicas stored in the metaspace.*
<b>Response Time</b>	The average response time for the metaspace.*
<b>Gets</b>	The total number of "get" operations performed on the user-spaces defined on the metaspace.*
<b>Gets/interval</b>	The number of "get" operations performed on the user-spaces defined for the metaspace during the current polling interval.*
<b>Gets/sec</b>	The rate of "get" operations (per second) performed on the user-spaces defined for the metaspace.*
<b>Puts</b>	The total number of "put" operations performed on the user-spaces defined on the metaspace.*
<b>Puts/interval</b>	The number of "put" operations performed on the user-spaces defined for the metaspace during the current polling interval.*
<b>Puts/sec</b>	The rate of "put" operations (per second) performed on the user-spaces defined for the metaspace.*
<b>Takes</b>	The total number of "take" operations performed on the user-spaces defined on the metaspace.*
<b>Takes/interval</b>	The number of "take" operations performed on the user-spaces defined for the metaspace during the current polling interval.*
<b>Takes/sec</b>	The rate of "take" operations (per second) performed on the user-spaces defined for the metaspace.*
<b>Expires</b>	The total number of entries in the user-spaces defined on the metaspace that have expired.*
<b>Expires/interval</b>	The number of entries in the user-spaces defined for the metaspace that expired during the current polling interval.*
<b>Expires/sec</b>	The rate of entries in the user-spaces defined for the metaspace that expired (per second).*
<b>Evicts</b>	The total number of entries in the user-spaces defined on the metaspace that have been evicted.*
<b>Evicts/interval</b>	The number of entries performed in the user-spaces defined for the metaspace that were evicted during the current polling interval.*
<b>Evicts/sec</b>	The rate of entries in the user-spaces defined for the metaspace that were evicted (per second).*
<b>Locks</b>	The total number of locks in the user-spaces defined for the metaspace.*

<b>Unlocks</b>	The total number of unlocks in the user-spaces defined for the metaspace.*
<b>Invokes</b>	The remote invocation count.*
<b>Queries</b>	The browser queries count in the user-spaces defined for the metaspace.*
<b>Misses</b>	The total number of misses on the user-spaces defined for the metaspace.*
<b>ToPersist</b>	The ToPersist count, which indicates how many tuples are required to be persisted to the database if the write-behind feature is configured.*
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > (Project Name) > <b>Solution Package Configuration</b> > <b>TIBCO Active Spaces</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Timestamp</b>	The date and time the row data was last updated.

## Metaspace Summary

This display provides a view of the current and historical metrics for a single metaspace. The trend graph in the bottom half of the display traces the current and historical total number of or rate data for gets, puts, takes, expires, and evictions, and also traces the average response time.






---

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the selected domain. Refer to TIBCO ActiveSpaces documentation for more information regarding these fields.

---

### Filter By:

The display might include these filtering options:

- Domain** Select the domain for which you want to show data in the display.
- Metaspace** Select the metaspace for which you want to show data in the display.

### Fields and Data:

- Last Update** The date and time in which the data in the display was last updated.
- Expired** When checked, performance data has not been received within the time specified (in seconds) in the **Expire Time** field in the **Duration** region in the RTView Configuration Application > (Project Name) > **Solution Package Configuration** > **TIBCO Active Spaces** > **DATA STORAGE** tab. The **Delete Time** field (also in the **Duration** region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
- Spaces** The number of user spaces defined in the metaspace.\*
- Members** The number of members (clients and servers) associated with the metaspace.\*
- Version** The metaspace's current version of TIBCO ActiveSpaces.
- Resp. Time** The average response time for the metaspace.\*
- Entries** The total number of entries stored in the metaspace.\*
- Replicas** The total number of replicas stored in the metaspace.\*

### Stats for last interval

- Gets**
  - Raw Counters--** The total number of gets for the metaspace.
  - for interval--** The number of gets for the current interval.
  - #/s --** The number of gets per second.
- Puts**
  - Raw Counters--** The total number of puts for the metaspace.
  - for interval--** The number of puts for the current interval.
  - #/s --** The number of puts per second.

<b>Takes</b>	<b>Raw Counters</b> -- The total number of takes for the metaspace. <b>for interval</b> -- The number of takes for the current interval. <b>#/s</b> -- The number of takes per second.
<b>Expirations</b>	<b>Raw Counters</b> -- The total number of expirations for the metaspace. <b>for interval</b> -- The number of expirations for the current interval. <b>#/s</b> -- The number of expirations per second.
<b>Evictions</b>	<b>Raw Counters</b> -- The total number of evictions for the metaspace. <b>for interval</b> -- The number of evictions for the current interval. <b>#/s</b> -- The number of evictions per second.
<b>Locks</b>	The total number of locks in the user-spaces defined for the metaspace.*
<b>Unlocks</b>	The total number of unlocks in the user-spaces defined for the metaspace.*
<b>Invokes</b>	The remote invocation count.*
<b>Queries</b>	The browser queries count in the user-spaces defined for the metaspace.*
<b>Misses</b>	The total number of misses in the user-spaces defined for the metaspace.*
<b>ToPersist</b>	The ToPersist count, which indicates how many tuples are required to be persisted to the database if the write-behind feature is configured.*

**Trends**

Traces the following:

**Gets(/s)** -- traces the total number of gets, or the number of gets per second with **Use Rates** selected.

**Puts(/s)**-- traces the total number of puts, or the number of puts per second with **Use Rates** selected.

**Takes(/s)** -- traces the total number of takes, or the number of takes per second with **Use Rates** selected.

**Expires(/s)** -- traces the total number of expires, or the number of expires per second with **Use Rates** selected.

**Evicts(/s)** -- traces the total number of evicts, or the number of evicts per second with **Use Rates** selected.

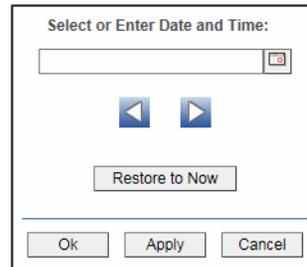
**Response Time** -- traces the average response time.

**Use Rates** Select this check box to trace the rates (**Gets/s, Puts/s, Takes/s, Expires/s, Evicts/s**) instead of the total numbers (**Gets, Puts, Takes, Expires, Evicts**).

**Log Scale** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.

**Base at Zero** Select to use zero (**0**) as the Y axis minimum for all graph traces.

**Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## All Spaces Table

The table in this display provides a view of all of your spaces and their associated metric data including domain, metaspace, space, alert level, alert count, and the current value of each gathered metric. You can click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by clicking a row to view details for the selected adapter in the ["Space Summary"](#) display.

ActiveSpaces All Spaces - Table 21-Apr-2017 15:55 Data OK

Domain: Production Metaspace: retailV214

Space Count: 3

Domain	Metaspace	Space	Alert Level	Alert Count	Space State	Members	Seeders
Production	retailV214	customers		0	READY	4	
Production	retailV214	inventory		0	READY	4	
Production	retailV214	stores		0	READY	4	

**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu , Table open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the selected domain. Refer to TIBCO ActiveSpaces documentation for more information regarding these fields.

**Filter By:**

- Domain** Select the domain for which you want to view data.
- Metaspace** Select the metaspace for which you want to view data.

**Spaces Table:**

- Domain** The name of the domain.
- Metaspace** The name of the metaspace.
- Space** The name of the space.

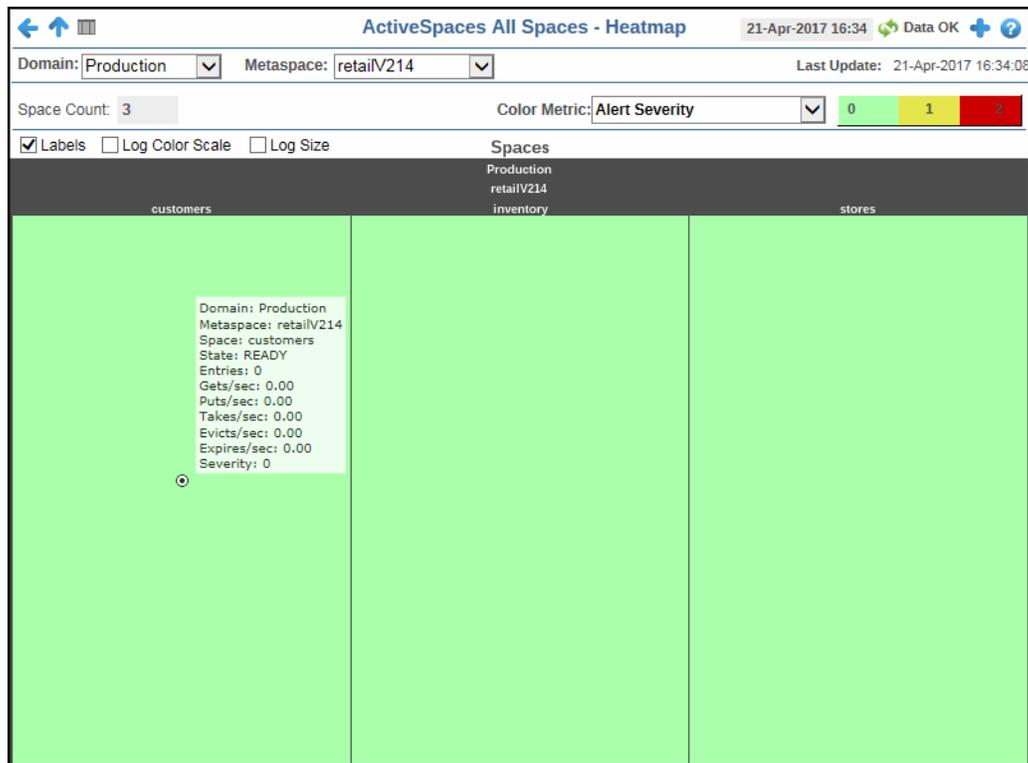
<b>Alert Level</b>	The current alert severity.  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
<b>Alert Count</b>	The total number of alerts for the host.
<b>Space State</b>	The current state of the space.*
<b>Members</b>	The total number of members in the space.*
<b>Seeders</b>	The number of seeders in the space.*
<b>Min Seeder Count</b>	The defined minimum seeder count (minimum number of seeders that need to be joined to the space before the space becomes ready).*
<b>CapacityPerSeeder</b>	The capacity value for the space in number of entries per seeder.*
<b>Entries</b>	The total number of entries stored in the space.*
<b>Replicas</b>	The total number of replicas stored in the space.*
<b>Gets</b>	The total number of "get" operations performed on the user-spaces defined on the space.*
<b>Gets/interval</b>	The number of "get" operations performed on the user-spaces defined for the space during the current polling interval.*
<b>Gets/sec</b>	The rate of "get" operations (per second) performed on the user-spaces defined for the space.*
<b>Puts</b>	The total number of "put" operations performed on the user-spaces defined on the space.*
<b>Puts/interval</b>	The number of "put" operations performed on the user-spaces defined for the space during the current polling interval.*
<b>Puts/sec</b>	The rate of "put" operations (per second) performed on the user-spaces defined for the space.*
<b>Takes</b>	The total number of "take" operations performed on the user-spaces defined on the space.*
<b>Takes/interval</b>	The number of "take" operations performed on the user-spaces defined for the space during the current polling interval.*
<b>Takes/sec</b>	The rate of "take" operations (per second) performed on the user-spaces defined for the space.*
<b>Expires</b>	The total number of entries in the user-spaces defined on the space that have expired.*
<b>Expires/interval</b>	The number of entries in the user-spaces defined for the space that expired during the current polling interval.*
<b>Expires/sec</b>	The rate of entries in the user-spaces defined for the space that expired (per second).*
<b>Evicts</b>	The total number of entries in the user-spaces defined on the space that have been evicted.*
<b>Evicts/interval</b>	The number of entries performed on the user-spaces defined for the space that were evicted during the current polling interval.*
<b>Evicts/sec</b>	The rate of entries in the user-spaces defined for the space that were evicted (per second).*

<b>Locks</b>	The total number of locks in the user-spaces defined for the space.*
<b>Unlocks</b>	The total number of unlocks in the user-spaces defined for the space.*
<b>Invokes</b>	The remote invocation count.*
<b>Queries</b>	The total number of queries in the user-spaces defined for the space.*
<b>Misses</b>	The total number of misses in the user-spaces defined for the space.*
<b>ToPersist</b>	The ToPersist count, which indicates how many tuples are required to be persisted to the database if the write-behind feature is configured.*
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > (Project Name) > <b>Solution Package Configuration</b> > <b>TIBCO Active Spaces</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Timestamp</b>	The date and time the row data was last updated.

## All Spaces Heatmap

This heatmap display provides an easy-to-view interface that allows you to quickly identify the current status of each of your spaces for each available metric. You can view the spaces in the heatmap based on the following metrics: current alert severity, entries, gets per second, puts per second, takes per second, expires per second, and evicts per second. By default, this display shows the heatmap based on the **Alert Severity** metric.

You can use the **Labels** check-box  to include or exclude labels in the heatmap, and you can mouse over a rectangle to see additional metrics for a space. Clicking one of the rectangles in the heatmap opens the "Space Summary" display, which allows you to see additional details for the selected space.



**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Filter By:**

- Domain** Select the domain for which you want to see data.
- Metaspace** Select the metaspace for which you want to see data.

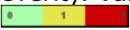
**Fields and Data:**

- Last Update** The date and time in which the data in the display was last updated.
- Space Count** The total number of spaces found for the selected Domain/Metaspace combination.
- Labels** Select this check box to display the names of the adapters at the top of each rectangle in the heatmap.
- Log Color Scale** Select this check box to use a logarithmic scale, rather than a linear scale, to map from the selected metric value for a cell to the color for the cell. **Log Scale** provides another way to distribute and differentiate values that you might not be able to see on a linear scale due to the dominant nature of large values in a linear scale.

**Log Size** Select this check box to use a logarithmic scale, rather than a linear scale, to map from the selected metric value for a cell to the size for the cell. **Log Scale** provides another way to distribute and differentiate values that you might not be able to see on a linear scale due to the dominant nature of large values in a linear scale.

**Color Metric** Choose a metric to view in the display.

**Alert Severity**

The current alert severity. Values range from **0** - **2**, as indicated in the color gradient  bar, where **2** is the highest Alert Severity:

-  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
-  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
-  Green indicates that no metrics have exceeded their alert thresholds.

**Entries**

The total number of entries in the space. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasSpaceEntriesHigh**. The middle value in the gradient bar indicates the middle value of the range.

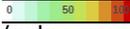
**Gets/sec**

The number of gets per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasSpaceGetRateHigh**. The middle value in the gradient bar indicates the middle value of the range.

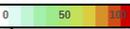
**Puts/sec**

The number of message sent per second. The color gradient bar  shows the range of the value/color mapping, ated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasSpacePutRateHigh**. The middle value in the gradient bar indicates the middle value of the range.

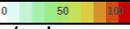
**Takes/sec**

The number of takes per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasSpaceTakeRateHigh**. The middle value in the gradient bar indicates the middle value of the range.

**Expires/sec**

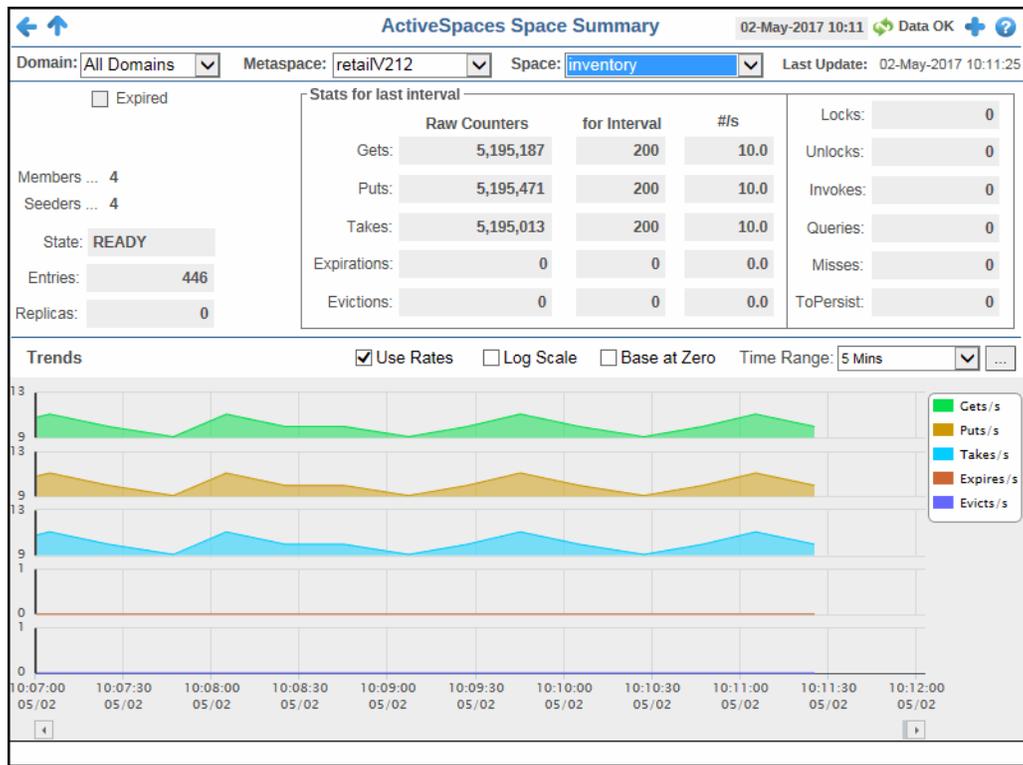
The number of expires per second. The color gradient bar  shows the range of the value/color mapping, ated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasSpaceExpireRateHigh**. The middle value in the gradient bar indicates the middle value of the range.

**Evicts/sec**

The number of evictions per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasSpaceEvictsRateHigh**. The middle value in the gradient bar indicates the middle value of the range.

## Space Summary

This display provides a view of the current and historical metrics for a single space. The trend graph in the bottom half of the display traces the current and historical total number of or rate data for gets, puts, takes, expires, and evictions.



**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the selected domain. Refer to TIBCO ActiveSpaces documentation for more information regarding these fields.

**Filter By:**

The display might include these filtering options:

- Domain** Select the domain for which you want to show data in the display.
- Metaspace** Select the metaspace for which you want to show data in the display.
- Space** Select the space for which you want to show data in the display.

**Fields and Data:**

<b>Last Update</b>	The date and time in which the data in the display was last updated.
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > (Project Name) > <b>Solution Package Configuration</b> > <b>TIBCO Active Spaces</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Members</b>	The total number of members associated with the space.* <b>Note:</b> You can click on this field to open the <a href="#">"Members by Space Table"</a> .
<b>Seeders</b>	The number of seeders in the space.* <b>Note:</b> You can click on this field to open the <a href="#">"Members by Space Table"</a> .
<b>State</b>	The current state of the space.*
<b>Entries</b>	The total number of entries stored in the space.*
<b>Replicas</b>	The total number of replicated entries in the space.*
<b>Stats for last interval</b>	
<b>Gets</b>	<b>Raw Counters</b> -- The total number of gets for the space. <b>for interval</b> -- The number of gets for the current interval. <b>#/s</b> -- The number of gets received per second.
<b>Puts</b>	<b>Raw Counters</b> -- The total number of puts for the space. <b>for interval</b> -- The number of puts for the current interval. <b>#/s</b> -- The number of puts received per second.
<b>Takes</b>	<b>Raw Counters</b> -- The total number of takes for the space. <b>for interval</b> -- The number of takes for the current interval. <b>#/s</b> -- The number of takes received per second.
<b>Expirations</b>	<b>Raw Counters</b> -- The total number of expirations for the space. <b>for interval</b> -- The number of expirations for the current interval. <b>#/s</b> -- The number of expirations received per second.
<b>Evictions</b>	<b>Raw Counters</b> -- The total number of evictions for the space. <b>for interval</b> -- The number of evictions for the current interval. <b>#/s</b> -- The number of evictions received per second.
<b>Locks</b>	The total number of locks in the user-spaces defined for the space.*
<b>Unlocks</b>	The total number of unlocks in the user-spaces defined for the space.*
<b>Invokes</b>	The remote invocation count.*
<b>Queries</b>	The total number of queries in the user-spaces defined for the space.*
<b>Misses</b>	The total number of misses on the user-spaces defined for the space.*
<b>ToPersist</b>	The ToPersist count, which indicates how many tuples are required to be persisted to the database if the write-behind feature is configured.*

**Trends**

Traces the following:

**Gets(/s)** -- traces the total number of gets, or the number of gets per second with **Use Rates** selected.

**Puts(/s)**-- traces the total number of puts, or the number of puts per second with **Use Rates** selected.

**Takes(/s)** -- traces the total number of takes, or the number of takes per second with **Use Rates** selected.

**Expires(/s)** -- traces the total number of expires, or the number of expires per second with **Use Rates** selected.

**Evicts(/s)** -- traces the total number of evicts, or the number of evicts per second with **Use Rates** selected.

**Response Time** -- traces the average response time.

**Use Rates** Select this check box to trace the rates (**Gets/s, Puts/s, Takes/s, Expires/s, Evicts/s**) instead of the total numbers (**Gets, Puts, Takes, Expires, Evicts**).

**Log Scale** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.

**Base at Zero** Select to use zero (**0**) as the Y axis minimum for all graph traces.

**Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar  .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## All Queries Table

This display allows you to view queries by domain, metaspace, and space and view the performance metrics for the queries. Clicking on a query in the table opens the “[Query Summary](#)” display.

Domain	Timestamp	Metaspace	space_name	QueryDuration	query_status	
tas_domain	16-Mar-2018 13:14:21	ms1	test_space_1	52	2	key/10 = 10
tas_domain	16-Mar-2018 13:14:21	ms1	test_space_1	20	2	value like "valu
tas_domain	16-Mar-2018 13:14:21	ms1	test_space_1	20	2	value like "valu
tas_domain	16-Mar-2018 13:14:21	ms1	test_space_1	20	2	value like "valu
tas_domain	16-Mar-2018 13:14:21	ms1	test_space_1	19	2	value like "valu
tas_domain	16-Mar-2018 13:14:21	ms1	test_space_1	18	2	key/10 = 10
tas_domain	16-Mar-2018 13:14:21	ms1	test_space_1	17	2	key/10 = 10
tas_domain	16-Mar-2018 13:14:21	ms1	test_space_1	16	2	key/10 = 10
tas_domain	16-Mar-2018 13:14:21	ms1	test_space_1	2	2	value like "valu
tas_domain	16-Mar-2018 13:14:21	ms1	test_space_1	2	2	value like "valu
tas_domain	16-Mar-2018 13:14:21	ms1	test_space_1	2	2	value like "valu
tas_domain	16-Mar-2018 13:14:21	ms1	test_space_1	2	2	value like "valu
tas_domain	16-Mar-2018 13:14:21	ms1	test_space_1	1	2	year(time) = 20
tas_domain	16-Mar-2018 13:14:21	ms1	test_space_1	1	2	year(time) = 20
tas_domain	16-Mar-2018 13:14:21	ms1	test_space_1	1	2	(key%10) in (0,
tas_domain	16-Mar-2018 13:14:21	ms1	test_space_1	1	2	value not like "\
tas_domain	16-Mar-2018 13:14:21	ms1	\$spacedef_space	0	2	
tas_domain	16-Mar-2018 13:14:21	ms1	\$space_state	0	2	
tas_domain	16-Mar-2018 13:14:21	ms1	\$members	0	2	
tas_domain	16-Mar-2018 13:14:21	ms1	\$space_members	0	2	space_name="
tas_domain	16-Mar-2018 13:14:21	ms1	\$members	0	2	
tas_domain	16-Mar-2018 13:14:21	ms1	\$members	0	2	
tas_domain	16-Mar-2018 13:14:21	ms1	\$spacedef_space	0	2	
tas_domain	16-Mar-2018 13:14:21	ms1	\$space_members	0	2	space_name="
tas_domain	16-Mar-2018 13:14:21	ms1	\$spacedef_space	0	2	
tas_domain	16-Mar-2018 13:14:21	ms1	\$space_state	0	2	

### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the selected domain. Refer to TIBCO ActiveSpaces documentation for more information regarding these fields.

### Filter By:

The display might include these filtering options:

- Domain** Select the domain for which you want to show data in the display.
- Metaspace** Select the metaspace for which you want to show data in the display.

**Space** Select the space for which you want to show data in the display.

#### Fields and Data:

**Query Count** The total number of queries listed in the table.

**Show Expired** Select this toggle to display expired queries in the table.

#### Queries Table

**Domain** The name of the domain containing the query.

**Timestamp** The date and time that the row in the table was last updated.

**Metaspace** The name of the metaspace containing the query.

**space\_name** The name of the space containing the query.

**Query Duration** The duration, in seconds, of the query.\*

**query\_status** The status of the query.\*

**0** - Failed

**1** - In progress

**2** - Completed

**filter** The filter used in the query.\*

**query\_type** The type of query.\*

**scan\_type** Lists whether the query used a table scan or an index scan.\*

**index\_name** The name of the index being used in the query.\*

**limit** Lists the maximum number of entries that can be returned when executing a query.\*

**estimated\_cost** The estimated execution time of the query.\*

**actual\_cost** The actual execution time of the query.\*

**abort** When checked, denotes that the query was aborted.\*

**StartTime** Start time of the query.

**EndTime** End time of the query.

**start\_time** Internal start time of the query.\*

**end\_time** Internal end time of the query.\*

**request\_id** The request id of the query.\*

**parent\_request\_id** The request id of the query's parent.\*

**member\_name** The name of the member node.\*

**member\_id** The id of the member node.\*

**process\_id** The process ID of the member node processing the query.\*

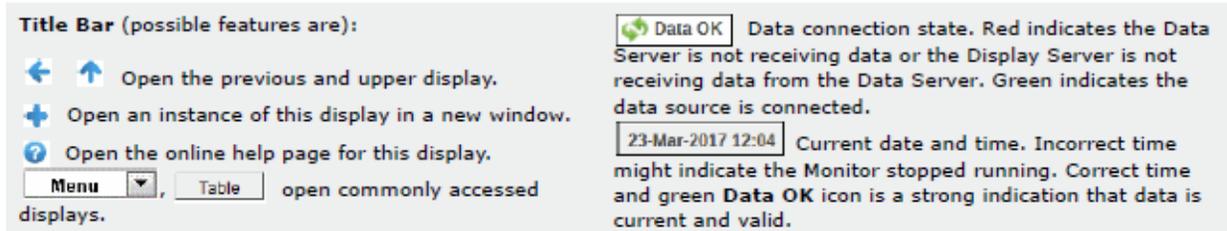
**Expired** When checked, performance data has not been received within the time specified (in seconds) in the **Expire Time** field in the **Duration** region in the RTView Configuration Application > (Project Name) > **Solution Package Configuration** > **TIBCO Active Spaces** > **DATA STORAGE** tab. The **Delete Time** field (also in the **Duration** region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.

## Query Summary

This display allows you to view performance metrics for a particular query, as well as to view any related queries. Data only appears in this display when you select a query from the "All Queries Table".

The screenshot displays the 'ActiveSpaces Query Summary' window. At the top, it shows navigation arrows, the title 'ActiveSpaces Query Summary', and a timestamp '16-Mar-2018 13:16' with a 'Data OK' indicator. Below the title are dropdown menus for 'Domain: tas\_domain', 'Metaspace: ms1', and 'Space: test\_space\_1'. The main content is divided into several sections:

- Selected Query:** A table with columns: Domain, Timestamp, Metaspace, space\_name, QueryDuration, query\_status. The row shows: tas\_domain, 16-Mar-2018 13:16:25, ms1, test\_space\_1, 52, 2 | key/10 = 10.
- Space Info:** Domain: tas\_domain, Metaspace: ms1, Space: test\_space\_1.
- Member Info ...:** Name: ms1\_as-agent-1, Id: c009c8d1-1b58-59d27f48-124, Process Id: 17664.
- Query Info:** Duration: 52, Status: 2, Start Time: Oct 4, 2017 5:33:58 PM, End Time: Oct 4, 2017 5:34:50 PM, Limit: 10000, Request Id: c009c8d1-1b58-59d27f48-124-37a, Query Type: Query, Scan Type: TableScan, Index name: KeyIndex, Estimated Cost: 10000000, Actual Cost: 10000000, Filter: key/10 = 10. There is an 'Abort' checkbox.
- Related Query Count:** 1. **Parent Request Id:** c009c8d1-1b58-59d57df5-132-48.
- Related Queries:** A table with columns: Domain, Timestamp, Metaspace, space\_name, QueryDuration, query\_status. The row shows: tas\_domain, 16-Mar-2018 13:16:25, ms1, test\_space\_1, 52, 2 | key/10 = 10.




---

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the selected domain. Refer to TIBCO ActiveSpaces documentation for more information regarding these fields.

---

### Filter By:

The display might include these filtering options:

- Domain** Select the domain for which you want to show data in the display.
- Metaspace** Select the metaspace for which you want to show data in the display.
- Space** Select the space for which you want to show data in the display.

### Fields and Data:

- Selected Query** Lists the details of the query selected from the ["All Queries Table"](#).

### Space Info

- Domain** The name of the domain in which the query resides.
- Metaspace** The name of the metaspace in which the query resides.
- Space** The name of the space in which the query resides.

### Member Info

**Note:** You can click this region to open the ["Member Summary"](#) display.

- Name** The name of the member node.\*
- Id** The id of the member node.\*
- Process ID** The process ID of the member node processing the query.\*

### Query Info

- Duration** The duration, in seconds, of the query.\*
- Status** The status of the query.\*
  - 0** - Failed
  - 1** - In progress
  - 2** - Completed
- Start Time** Start time of the query.
- End Time** End time of the query.

<b>Limit</b>	Lists the maximum number of entries that can be returned when executing a query.*
<b>Abort</b>	When checked, denotes that the query was aborted.*
<b>Request Id</b>	The request id of the query.*
<b>Query Type</b>	The type of query.*
<b>Scan Type</b>	Lists whether the query used a table scan or an index scan.*
<b>Index Name</b>	The name of the index being used in the query.*
<b>Estimated Cost</b>	The estimated execution time of the query.*
<b>Actual Cost</b>	The actual execution time of the query.*
<b>Filter</b>	The filter used in the query.*
<b>Related Query Count</b>	The number of queries related to the selected query.
<b>Parent Request Id</b>	The request ID of the query's parent.
<b>Related Queries</b>	Lists the details of any related ("sibling") queries.

## Members View

The displays in this view allow you to view the current and historical metrics for all members in a particular metaspace, view data for members within a particular space, and view data for all spaces for a particular member. The available displays in this View are:

- ["All Members Table"](#): A tabular view of all members in a particular metaspace.
- ["All Members Heatmap"](#): A heatmap view of all members in a particular metaspace.
- ["Member Summary"](#): This display allows you to view current and trending data for a single member for a particular metaspace.
- ["Member Summary - Process"](#): This display allows you to view current and trending process statistics for a single member for a particular metaspace.
- ["Member Summary -JVM"](#): This display allows you to view current and trending JVM statistics for a single member for a particular metaspace.
- ["Members by Space Table"](#): A tabular view of all members in a particular space.
- ["Members by Space Heatmap"](#): A heatmap view of all members in a particular space.
- ["Spaces by Member Table"](#): A tabular view of all spaces for a particular member.
- ["Member by Space Summary"](#): This display allows you to view data for a selected member for a particular space.

## All Members Table

The table in this display provides a view of all of the members in a particular metaspace and their associated metric data including domain, metaspace, alert severity, alert count, and the current value of each gathered metric. You can click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by clicking a row to view details for the selected member in the "Member Summary" display

Domain	Metaspace	Member Name	Alert Level	Alert Count	Management Role	Host Address	Host Po
Production	retailV213	AlphaTestBed	●	0	MEMBER	192.168.200.134	6
Production	retailV213	as-agent-1	●	0	MEMBER	192.168.200.131	6
Production	retailV213	retail_get	●	0	MEMBER	192.168.200.131	6
Production	retailV213	retail_put	●	0	MANAGER	192.168.200.131	6
Production	retailV213	retail_take	●	0	MEMBER	192.168.200.131	6
Production	retailV214	as-agent-1	●	0	MEMBER	192.168.200.74	6
Production	retailV214	retail_get	●	0	MEMBER	192.168.200.74	6
Production	retailV214	retail_put	●	0	MANAGER	192.168.200.74	6
Production	retailV214	retail_take	●	0	MEMBER	192.168.200.74	6
Production	retailV216	as-agent-1	●	0	MEMBER	192.168.200.71	6
Production	retailV216	as-agent-2	●	0	MEMBER	192.168.200.71	6
Production	retailV216	retail_get	●	0	MEMBER	192.168.200.71	6
Production	retailV216	retail_put	●	0	MANAGER	192.168.200.71	6
Production	retailV216	retail_take	●	0	MEMBER	192.168.200.71	6

**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the selected domain. Refer to TIBCO ActiveSpaces documentation for more information regarding these fields.

**Filter By:**

- Domain** Select the domain for which you want to view data.
- Metaspace** Select the metaspace for which you want to view data.

<b>Member Count</b>	The resulting total number of members found in the filtered query, and listed in the <b>Members</b> table.
<b>Members Table</b>	
<b>Domain</b>	The name of the domain.
<b>Metaspace</b>	The name of the metaspace.
<b>Member Name</b>	The name of the member.
<b>Alert Level</b>	The current alert severity.  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
<b>Alert Count</b>	The total number of alerts for the host.
<b>Management Role</b>	The member's role within the metaspace.
<b>Host Address</b>	The IP address of the host.*
<b>Host Port</b>	The port of the host.*
<b>ProcessID</b>	The process ID of the process being monitored.*
<b>Process Name</b>	The name of the process.*
<b>NumSpaces</b>	The number of spaces in which the metaspace member is a member.*
<b>Entries</b>	The number of entries associated with the member.*
<b>Replicas</b>	The number of replicas.*
<b>Gets</b>	The total number of "get" operations performed on the user-spaces defined on the metaspace.*
<b>Gets/interval</b>	The number of "get" operations performed on the user-spaces defined for the metaspace during the current polling interval.*
<b>Gets/sec</b>	The rate of "get" operations (per second) performed on the user-spaces defined for the metaspace.*
<b>Puts</b>	The total number of "put" operations performed on the user-spaces defined on the metaspace.*
<b>Puts/interval</b>	The number of "put" operations performed on the user-spaces defined for the metaspace during the current polling interval.*
<b>Puts/sec</b>	The rate of "put" operations (per second) performed on the user-spaces defined for the metaspace.*
<b>Takes</b>	The total number of "take" operations performed on the user-spaces defined on the metaspace.*
<b>Takes/interval</b>	The number of "take" operations performed on the user-spaces defined for the metaspace during the current polling interval.*
<b>Takes/sec</b>	The rate of "take" operations (per second) performed on the user-spaces defined for the metaspace.*
<b>Expires</b>	The total number of entries in the user-spaces defined on the metaspace that have expired.*

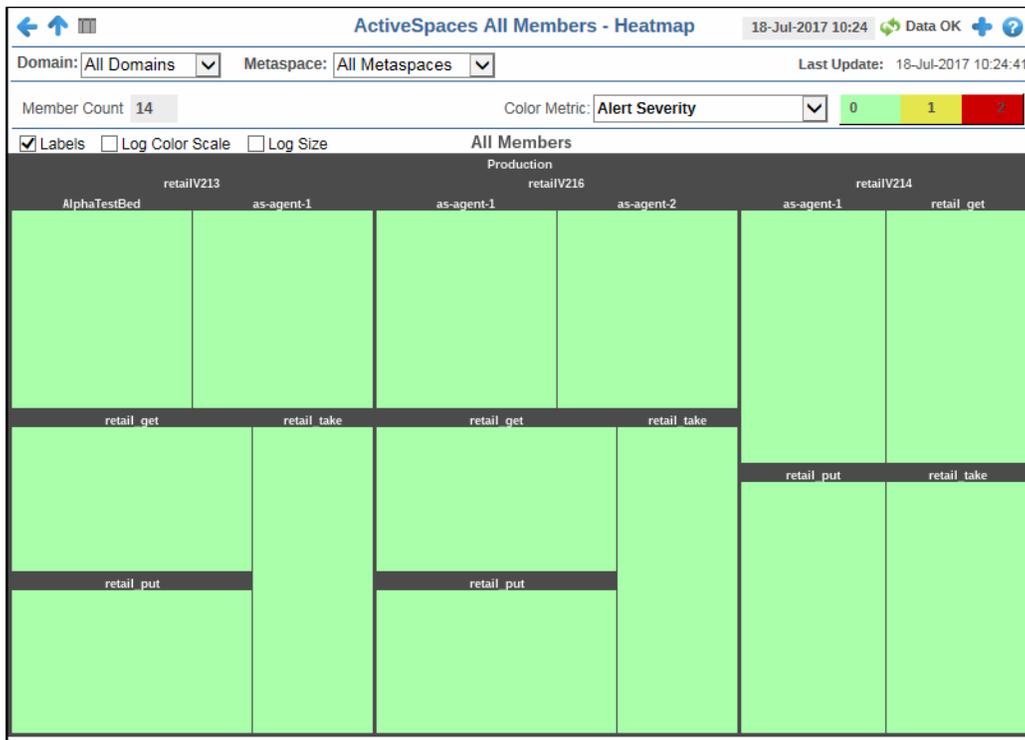
<b>Expires/interval</b>	The number of entries performed in the user-spaces defined for the metaspace that expired during the current polling interval.*
<b>Expires/sec</b>	The rate of entries in the user-spaces defined for the metaspace that expired (per second).*
<b>Evicts</b>	The total number of entries in the user-spaces defined on the metaspace that have been evicted.*
<b>Evicts/interval</b>	The number of entries performed in the user-spaces defined for the metaspace that were evicted during the current polling interval.*
<b>Evicts/sec</b>	The rate of entries in the user-spaces defined for the metaspace that were evicted (per second).*
<b>Locks</b>	The total number of locks in the user-spaces defined for the metaspace.*
<b>Unlocks</b>	The total number of unlocks in the user-spaces defined for the metaspace.*
<b>Invokes</b>	The remote invocation count.*
<b>Queries</b>	The total number of queries in the user-spaces defined for the metaspace.*
<b>Misses</b>	The total number of misses in the user-spaces defined for the metaspace.*
<b>ToPersist</b>	The ToPersist count, which indicates how many tuples are required to be persisted to the database if the write-behind feature is configured.*
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > (Project Name) > <b>Solution Package Configuration</b> > <b>TIBCO Active Spaces</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Join Time</b>	The time that the member joined the space.
<b>Timestamp</b>	The date and time the row data was last updated.
<b>sys_name</b>	The operating system on which the member is running.*
<b>cmd_name</b>	Indicates the command used to start <b>as-admin</b> .*
<b>user_name</b>	The name of the user running the process.*
<b>thread_count</b>	The number of threads running for the process.*
<b>res_mem_size</b>	Indicates the amount of physical memory currently allocated to the member.*
<b>mem_load</b>	The percentage of memory being used.*
<b>peak_res_mem_size</b>	Indicates the peak size of the system resident memory allocated by the system.*
<b>page_size</b>	Indicates the current size of the system pagefiles allocated by the system.*
<b>peak_page_size</b>	Indicates the maximum size of the system pagefiles allowed by the system.*
<b>process_cpu_load</b>	Indicates the load on the CPU (CPU percentage).*
<b>cpu_count</b>	The number of CPUs running on the system.*
<b>jvm_comm_heap_size</b>	The committed JVM heap usage, in megabytes.*

<b>jvm_max_heap_size</b>	The maximum JVM heap usage, in megabytes.*
<b>jvm_used_heap_size</b>	The used JVM heap, in megabytes.*
<b>jvm_comm_nonheap_size</b>	The committed JVM non-heap memory usage, in megabytes.*
<b>jvm_max_nonheap_size</b>	The maximum JVM non-heap memory usage, in megabytes.*
<b>jvm_used_nonheap_size</b>	The used JVM non-heap memory, in megabytes.*
<b>jvm_finalizing_count</b>	The amount of memory freed by the finalize operation on the JVM.*
<b>as_version</b>	The current ActiveSpaces version running.*
<b>JVMMemoryUsedPercent</b>	The percentage of memory used by the JVM.*

## All Members Heatmap

This heatmap display provides an easy-to-view interface that allows you to quickly identify the current status of each of your members for each available metric. You can view the members in the heatmap based on the following metrics: the current alert severity, the number of entries, the number of gets per second, the number of puts per second, the number of takes per second, the number of expires per second, the number of evictions per second, the percentage of CPU used, the percentage of memory used, and the percentage of JVM memory used. By default, this display shows the heatmap based on the **Alert Severity** metric.

You can use the **Labels** check-box  to include or exclude labels in the heatmap, and you can mouse over a rectangle to see additional metrics for a particular member. Clicking one of the rectangles in the heatmap opens the "[Member Summary](#)" display, which allows you to see additional details for the selected member.



**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Filter By:**

**Domain** Select the Domain for which you want to view data.

**Metaspace** Select the metaspace for which you want to view data.

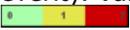
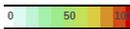
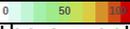
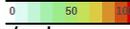
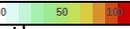
**Fields and Data:**

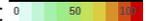
**Last Update** The date and time in which the data in the display was last updated.

**Member Count** The number of members found for the selected **Domain/Metaspace** combination.

**Labels** Select this check box to display the names of the adapters at the top of each rectangle in the heatmap.

**Log Color Scale** Select this check box to use a logarithmic scale, rather than a linear scale, to map from the selected metric value for a cell to the color for the cell. **Log Scale** provides another way to distribute and differentiate values that you might not be able to see on a linear scale due to the dominant nature of large values in a linear scale.

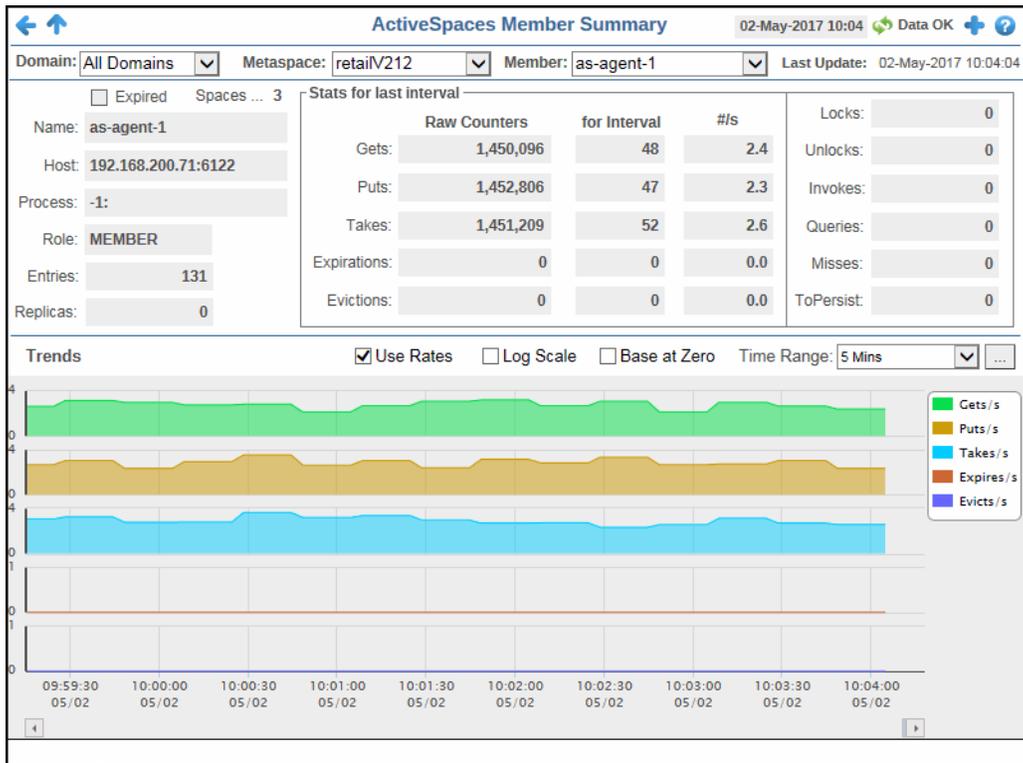
- Log Size** Select this check box to use a logarithmic scale, rather than a linear scale, to map from the selected metric value for a cell to the size for the cell. **Log Scale** provides another way to distribute and differentiate values that you might not be able to see on a linear scale due to the dominant nature of large values in a linear scale.
- Color Metric** Choose a metric to view in the display.
- Alert Severity** The current alert severity. Values range from **0** - **2**, as indicated in the color gradient  bar, where **2** is the highest Alert Severity:
-  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
  -  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
  -  Green indicates that no metrics have exceeded their alert thresholds.
- Entries** The total number of entries in the adapters. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasMemberEntriesHigh**. The middle value in the gradient bar indicates the middle value of the range.
- Gets/sec** The number of gets per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasMemberGetRateHigh**. The middle value in the gradient bar indicates the middle value of the range.
- Puts/sec** The number of puts per second. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasMemberPutRateHigh**. The middle value in the gradient bar indicates the middle value of the range.
- Takes/sec** The number of takes per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasMemberTakeRateHigh**. The middle value in the gradient bar indicates the middle value of the range.
- Expires/sec** The number of expires per second. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasMemberExpireRateHigh**. The middle value in the gradient bar indicates the middle value of the range.
- Evicts/sec** The number of evictions per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasMemberEvictsRateHigh**. The middle value in the gradient bar indicates the middle value of the range.
- CPU %** The percentage of CPU used. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasMemberCpuHigh**. The middle value in the gradient bar indicates the middle value of the range.

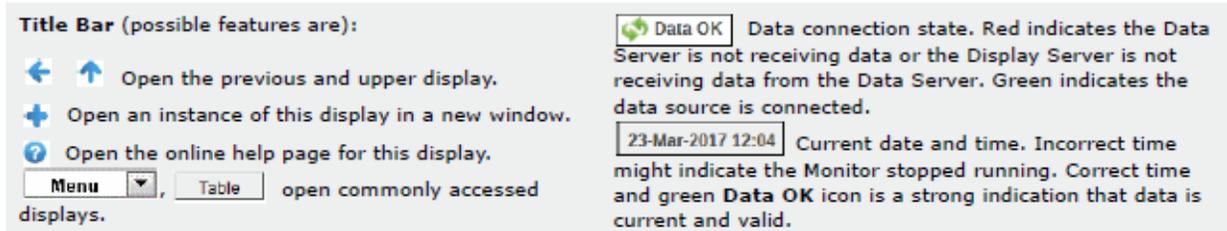
**Memory %** The percentage of memory used. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasMemberMemoryUsedHigh**. The middle value in the gradient bar indicates the middle value of the range.

**JVM Memory %** The percentage of JVM memory used. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasMemberJvmMemoryUsedHigh**. The middle value in the gradient bar indicates the middle value of the range.

## Member Summary

This display provides a view of the current and historical metrics for a single member. The trend graph in the bottom half of the display traces the current and historical total number of or rate data for gets, puts, takes, expires, and evictions.






---

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the selected domain. Refer to TIBCO ActiveSpaces documentation for more information regarding these fields.

---

### Filter By:

The display might include these filtering options:

- Domain** Select the domain for which you want to show data in the display.
- Metaspace** Select the metaspace for which you want to show data in the display.
- Member** Select the space for which you want to show data in the display.

### Fields and Data:

- Last Update** The date and time in which the data in the display was last updated.
- Expired** When checked, performance data has not been received within the time specified (in seconds) in the **Expire Time** field in the **Duration** region in the RTView Configuration Application > (Project Name) > **Solution Package Configuration** > **TIBCO Active Spaces** > **DATA STORAGE** tab. The **Delete Time** field (also in the **Duration** region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
- Spaces** The total number of spaces in which the member is a member.\*  
**Note:** Clicking on this field opens the "[Spaces by Member Table](#)" display.
- Name** The name of the member.
- Host** The IP address of the host.
- Process** The process ID and process name (ProcessID:ProcessName).\*
- Role** The role of the member.
- Entries** The number of entries for the member.\*
- Replicas** The number of replicas for the member.\*
- Stats for last interval**

<b>Gets</b>	<p><b>Raw Counters</b>-- The total number of "get" operations performed on the user-spaces defined on the metaspace.*</p> <p><b>for Interval</b>-- The number of "get" operations performed on the user-spaces defined for the metaspace during the current polling interval.*</p> <p><b>#/s</b> -- The rate of "get" operations (per second) performed on the user-spaces defined for the metaspace.*</p>
<b>Puts</b>	<p><b>Raw Counters</b>-- The total number of "put" operations performed on the user-spaces defined on the metaspace.*</p> <p><b>for Interval</b>-- The number of "put" operations performed on the user-spaces defined for the metaspace during the current polling interval.*</p> <p><b>#/s</b> -- The rate of "put" operations (per second) performed on the user-spaces defined for the metaspace.*</p>
<b>Takes</b>	<p><b>Raw Counters</b>-- The total number of "take" operations performed on the user-spaces defined on the metaspace.*</p> <p><b>for Interval</b>-- The number of "take" operations performed on the user-spaces defined for the metaspace during the current polling interval.*</p> <p><b>#/s</b> -- The rate of "take" operations (per second) performed on the user-spaces defined for the metaspace.*</p>
<b>Expirations</b>	<p><b>Raw Counters</b>-- The total number of entries in the user-spaces defined on the metaspace that have expired.*</p> <p><b>for Interval</b>-- The number of entries performed in the user-spaces defined for the metaspace that expired during the current polling interval.*</p> <p><b>#/s</b> -- The rate of entries in the user-spaces defined for the metaspace that expired (per second).*</p>
<b>Evictions</b>	<p><b>Raw Counters</b>-- The total number of entries in the user-spaces defined on the metaspace that have been evicted.*</p> <p><b>for Interval</b>-- The number of entries performed in the user-spaces defined for the metaspace that were evicted during the current polling interval.*</p> <p><b>#/s</b> -- The rate of entries in the user-spaces defined for the metaspace that were evicted (per second).*</p>
<b>Locks</b>	The total number of locks in the user-spaces defined for the metaspace.*
<b>Unlocks</b>	The total number of unlocks in the user-spaces defined for the metaspace.*
<b>Invokes</b>	The remote invocation count.*
<b>Queries</b>	The total number of queries in the user-spaces defined for the metaspace.*
<b>Misses</b>	The total number of misses in the user-spaces defined for the metaspace.*
<b>ToPersist</b>	The ToPersist count, which indicates how many tuples are required to be persisted to the database if the write-behind feature is configured.*

**Trends**

Traces the following:

**Gets(/s)** -- traces the total number of gets, or the number of gets per second with **Use Rates** selected.

**Puts(/s)**-- traces the total number of puts, or the number of puts per second with **Use Rates** selected.

**Takes(/s)** -- traces the total number of takes, or the number of takes per second with **Use Rates** selected.

**Expires(/s)** -- traces the total number of expires, or the number of expires per second with **Use Rates** selected.

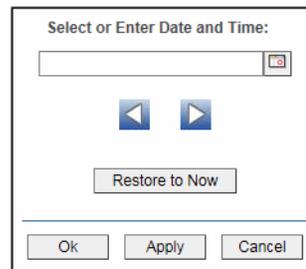
**Evicts(/s)** -- traces the total number of evicts, or the number of evicts per second with **Use Rates** selected.

**Use Rates** Select this check box to trace the rates (**Gets/s, Puts/s, Takes/s, Expires/s, Evicts/s**) instead of the total numbers (**Gets, Puts, Takes, Expires, Evicts**).

**Log Scale** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.

**Base at Zero** Select to use zero (**0**) as the Y axis minimum for all graph traces.

**Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

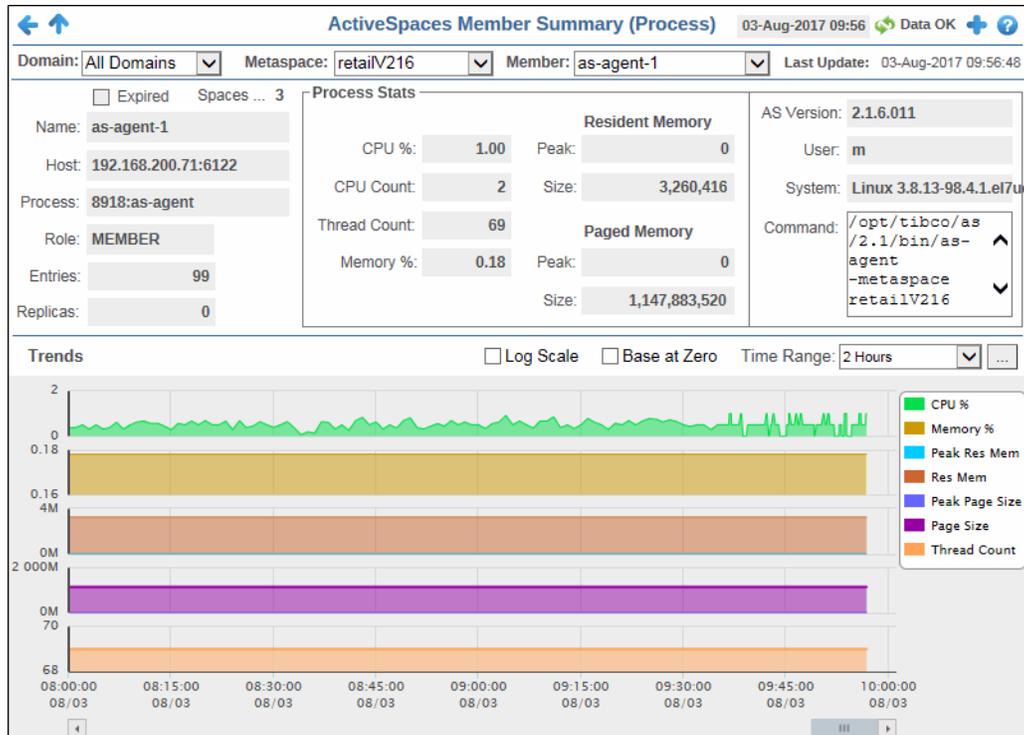
Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Member Summary - Process

This display provides a view of the current and historical process metrics for a single member. The trend graph in the bottom half of the display traces the current and historical process statistics for the selected member.

**Note:** These metrics are only available for members where system monitoring is enabled.



**Title Bar (possible features are):**

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu, Table open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the selected domain. Refer to TIBCO ActiveSpaces documentation for more information regarding these fields.

**Filter By:**

The display might include these filtering options:

<b>Domain</b>	Select the domain for which you want to show data in the display.
<b>Metaspace</b>	Select the metaspace for which you want to show data in the display.
<b>Member</b>	Select the space for which you want to show data in the display.

**Fields and Data:**

<b>Last Update</b>	The date and time in which the data in the display was last updated.
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > (Project Name) > <b>Solution Package Configuration</b> > <b>TIBCO Active Spaces</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Spaces</b>	The total number of spaces in which the member is a member. <b>Note:</b> Clicking on this field opens the " <a href="#">Spaces by Member Table</a> " display.
<b>Name</b>	The name of the member.
<b>Host</b>	The IP address of the host.*
<b>Process</b>	The process ID and process name (ProcessID:ProcessName).*
<b>Role</b>	The role of the member.
<b>Entries</b>	The number of entries for the member.*
<b>Replicas</b>	The number of replicas for the member.*
<b>Process Stats</b>	
<b>CPU %</b>	Indicates the load on the CPU (CPU percentage).*
<b>CPU Count</b>	The number of CPUs running on the system.*
<b>Thread Count</b>	The number of threads running for the process.*
<b>Memory %</b>	The percentage of memory being used.*
<b>Resident Memory</b>	<b>Peak</b> -- Indicates the peak size of the system resident memory allocated by the system.* <b>Size</b> -- Indicates the amount of physical memory currently allocated to the member.*
<b>Paged Memory</b>	<b>Peak</b> -- Indicates the maximum size of the system pagefiles allowed by the system.* <b>Size</b> -- Indicates the current size of the system pagefiles allocated by the system.*
<b>AS Version</b>	The current ActiveSpaces version running.*
<b>User</b>	The name of the user running the process.*
<b>System</b>	The operating system on which the member is running.*

**Command** Indicates the command used to start the member process.\*

## Trends

Traces the following:

**CPU %**-- traces the percentage of CPU being used.

**Memory %**-- traces the percentage of memory being used.\*

**Peak Res Mem**-- traces the peak size of the system resident memory allocated by the system.\*

**Res Mem**-- traces the amount of physical memory currently allocated to the member.\*

**Peak Page Size**-- traces the maximum size of the system pagefiles allowed by the system.\*

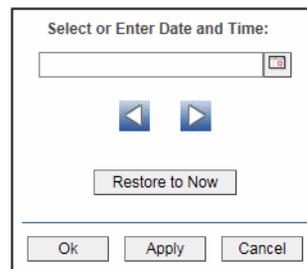
**Page Size**-- traces the current size of the system pagefiles allocated by the system.\*

**Thread Count**-- traces the number of threads running for the process.\*

**Log Scale** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.

**Base at Zero** Select to use zero (0) as the Y axis minimum for all graph traces.

**Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

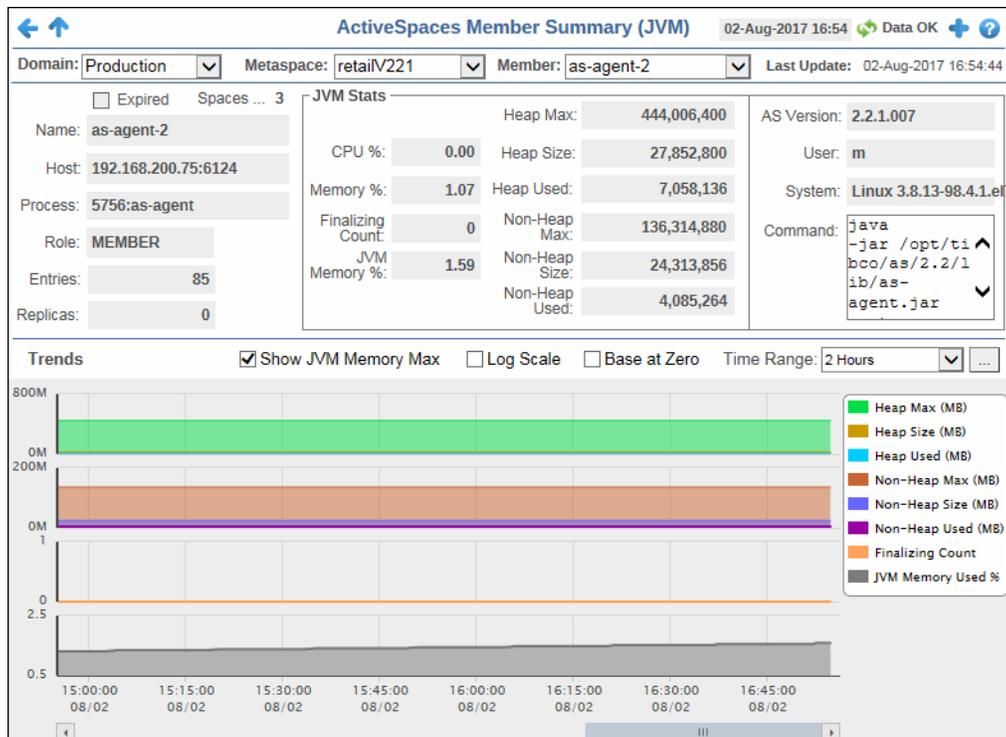
Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Member Summary -JVM

This display provides a view of the current and historical JVM statistics for a single member. The trend graph in the bottom half of the display traces the current and historical JVM metrics for the selected member.

**Note:** These metrics are only available for Java members where system monitoring is enabled.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the selected domain. Refer to TIBCO ActiveSpaces documentation for more information regarding these fields.

**Filter By:**

The display might include these filtering options:

<b>Domain</b>	Select the domain for which you want to show data in the display.
<b>Metaspace</b>	Select the metaspace for which you want to show data in the display.
<b>Member</b>	Select the space for which you want to show data in the display.

**Fields and Data:**

<b>Last Update</b>	The date and time in which the data in the display was last updated.
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > (Project Name) > <b>Solution Package Configuration</b> > <b>TIBCO Active Spaces</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Spaces</b>	The total number of spaces in which the member is a member.* <b>Note:</b> Clicking on this field opens the " <a href="#">Spaces by Member Table</a> " display.
<b>Name</b>	The name of the member.
<b>Host</b>	The IP address of the host.*
<b>Process</b>	The process ID and process name (ProcessID:ProcessName).*
<b>Role</b>	The role of the member.
<b>Entries</b>	The number of entries for the member.*
<b>Replicas</b>	The number of replicas for the member.*
<b>JVM Stats</b>	
<b>CPU %</b>	The load on the CPU (CPU percentage).*
<b>Memory %</b>	The percentage of memory being used.*
<b>Finalizing Count</b>	The amount of memory freed by the finalize operation on the JVM.*
<b>JVM Memory %</b>	The percentage of Java memory used by the JVM.*
<b>Heap Max</b>	The maximum JVM heap memory that can be used, in megabytes.*
<b>Heap Size</b>	The committed JVM heap size, in megabytes.*
<b>Heap Used</b>	The JVM heap memory currently being used, in megabytes.*
<b>Non-Heap Max</b>	The maximum JVM non-heap memory that can be used, in megabytes.*
<b>Non-Heap Size</b>	The committed JVM non-heap size, in megabytes.*
<b>Non-Heap Used</b>	The JVM non-heap memory currently being used, in megabytes.*
<b>AS Version</b>	The current ActiveSpaces version running.*

<b>User</b>	The name of the user running the process.*
<b>System</b>	The operating system on which the member is running.*
<b>Command</b>	Indicates the commands used to start the member process.*

**Trends**

Traces the following:

**Heap Max (MB)**- traces the maximum JVM heap memory that can be used, in megabytes.\*

**Heap Size (MB)**-- traces the maximum JVM heap usage, in megabytes.\*

**Heap Used (MB)**-- traces the committed JVM heap size, in megabytes.\*

**Non-Heap Max (MB)**-- traces the maximum JVM non-heap memory that can be used, in megabytes.\*

**Non-Heap Size (MB)**-- traces the committed JVM non-heap size, in megabytes.\*

**Non-Heap Used (MB)**-- traces the JVM non-heap memory currently being used, in megabytes.\*

**Finalizing Count**-- traces the amount of memory freed by the finalize operation on the JVM.\*

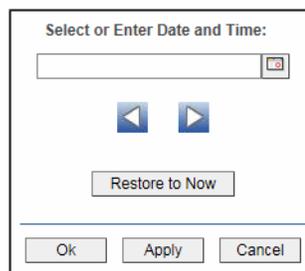
**JVM Memory Used %**-- traces the percentage of Java memory used by the JVM.\*

**Show JVM Memory Max** When selected, enables the Heap Max (MB) and Non-Heap Max (MB) metrics in the trend graph, which might be useful for removing the maximum metrics from the plot when they differ significantly from the used and committed values.

**Log Scale** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.

**Base at Zero** Select to use zero (0) as the Y axis minimum for all graph traces.

**Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Members by Space Table

The table in this display provides a view of all of your members and their associated metric data including domain, metaspace, space, alert severity, alert count, and the current value of each gathered metric. You can click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by clicking a row to view details for the selected member in the "Member by Space Summary" display.

Domain	Metaspace	Space	MemberName	Alert Severity	Alert Count	Drill Down
Production	retailV212	inventory	as-agent-1	0	0	SEE
Production	retailV212	inventory	retail_get	0	0	SEE
Production	retailV212	inventory	retail_put	0	0	SEE
Production	retailV212	inventory	retail_take	0	0	SEE

**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the selected domain. Refer to TIBCO ActiveSpaces documentation for more information regarding these fields.

**Filter By:**

<b>Domain</b>	Select the domain for which you want to view data.
<b>Metaspace</b>	Select the metaspace for which you want to view data.
<b>Space</b>	Select the space for which you want to view data.

**Member Count** The resulting total number of members found in the filtered query, which are listed in the **Members by Space** table.

**Members by Space Table**

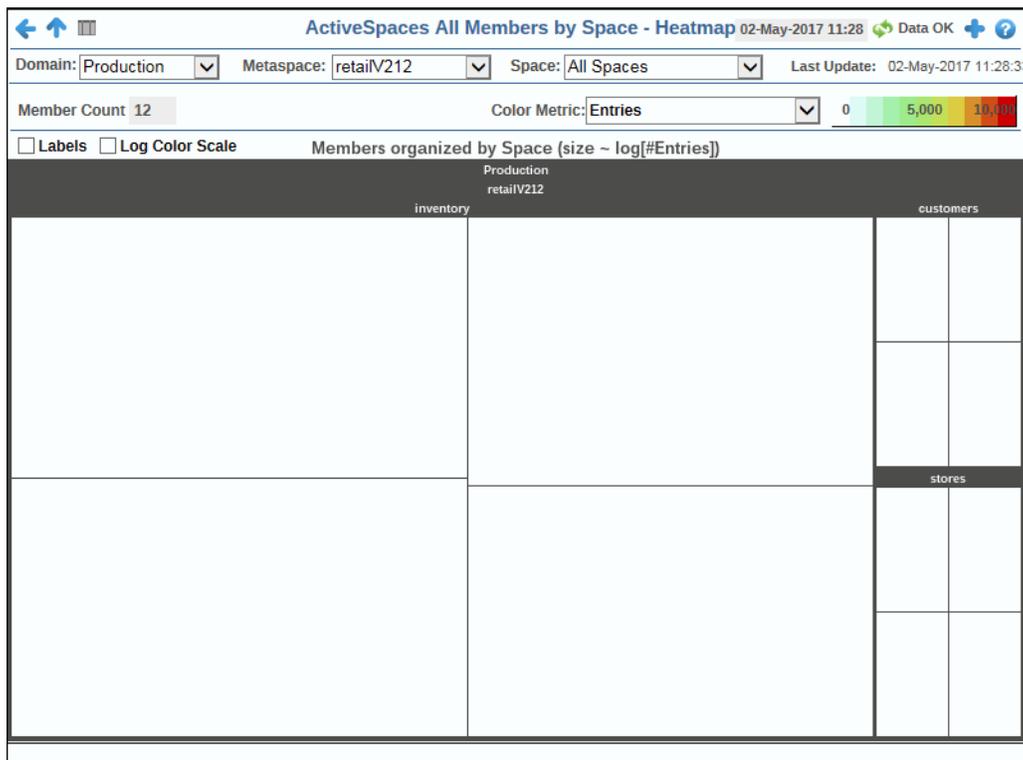
<b>Domain</b>	The name of the domain.
<b>Metaspace</b>	The name of the metaspace.
<b>Space</b>	The name of the space.
<b>Member Name</b>	The name of the member.
<b>Alert Severity</b>	The current alert severity.  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
<b>Alert Count</b>	The total number of alerts for the host.
<b>Distribution Role</b>	The member's role within the space.*
<b>Entries</b>	The number of entries.*
<b>% Capacity</b>	The percentage of available entries used for the space.
<b>Replicas</b>	The number of replicas.*
<b>Gets</b>	The total number of "get" operations performed on the user-spaces defined on the metaspace.*
<b>Gets/interval</b>	The number of "get" operations performed on the user-spaces defined for the metaspace during the current polling interval.*
<b>Gets/sec</b>	The rate of "get" operations (per second) performed on the user-spaces defined for the metaspace.*
<b>Puts</b>	The total number of "put" operations performed on the user-spaces defined on the metaspace.*
<b>Puts/interval</b>	The number of "put" operations performed on the user-spaces defined for the metaspace during the current polling interval.*
<b>Puts/sec</b>	The rate of "put" operations (per second) performed on the user-spaces defined for the metaspace.*
<b>Takes</b>	The total number of "take" operations performed on the user-spaces defined on the metaspace.*
<b>Takes/interval</b>	The number of "take" operations performed on the user-spaces defined for the metaspace during the current polling interval.*
<b>Takes/sec</b>	The rate of "take" operations (per second) performed on the user-spaces defined for the metaspace.*

<b>Expires</b>	The total number of entries in the user-spaces defined on the metaspace that have expired.*
<b>Expires/interval</b>	The number of entries performed in the user-spaces defined for the metaspace that expired during the current polling interval.*
<b>Expires/sec</b>	The rate of entries in the user-spaces defined for the metaspace that expired (per second).*
<b>Evicts</b>	The total number of entries in the user-spaces defined on the metaspace that have been evicted.*
<b>Evicts/interval</b>	The number of entries performed in the user-spaces defined for the metaspace that were evicted during the current polling interval.*
<b>Evicts/sec</b>	The rate of entries in the user-spaces defined for the metaspace that were evicted (per second).*
<b>Locks</b>	The total number of locks in the user-spaces defined for the metaspace.*
<b>Unlocks</b>	The total number of unlocks in the user-spaces defined for the metaspace.*
<b>Invokes</b>	The remote invocation count.*
<b>Queries</b>	The total number of queries in the user-spaces defined for the metaspace.*
<b>Misses</b>	The total number of misses in the user-spaces defined for the metaspace.*
<b>ToPersist</b>	The ToPersist count, which indicates how many tuples are required to be persisted to the database if the write-behind feature is configured.*
<b>ClientAvgGetMicros</b>	The client's average "get" latency in microseconds.*
<b>ClientAvgPutMicros</b>	The client's average "put" latency in microseconds.*
<b>ClientAvgTakeMicros</b>	The client's average "take" latency in microseconds.*
<b>ClientMaxGetMicros</b>	The client's highest "get" latency in microseconds.*
<b>ClientMaxPutMicros</b>	The client's highest "put" latency in microseconds.*
<b>ClientMaxTakeMicros</b>	The client's highest "take" latency in microseconds.*
<b>ClientTotalGetMillis</b>	The client's cumulative total "get" latency in milliseconds for the current polling period.*
<b>ClientTotalPutMillis</b>	The client's cumulative total "put" latency in milliseconds for the current polling period.*
<b>ClientTotalTakeMillis</b>	The client's cumulative total "take" latency in milliseconds for the current polling period.*
<b>ClientTotalMissMillis</b>	The client's cumulative total "miss" latency in milliseconds for the current polling period.*
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > (Project Name) > <b>Solution Package Configuration</b> > <b>TIBCO Active Spaces</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Timestamp</b>	The date and time the row data was last updated.

## Members by Space Heatmap

This heatmap display provides an easy-to-view interface that allows you to quickly identify the current status of each of your members for each available metric. You can view the members in the heatmap based on the following metrics: the number of entries, the number of gets per second, the number of puts per second, the number of takes per second, and the number of expires per second, and the number of evictions per second. By default, this display shows the heatmap based on the **Entries** metric.

You can use the **Labels** check-box  to include or exclude labels in the heatmap, and you can mouse over a rectangle to see additional metrics for a member. Clicking one of the rectangles in the heatmap opens the "Member Summary" display, which allows you to see additional details for the selected member.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

### Filter By:

- Domain** Select the Domain for which you want to view data.
- Metaspace** Select the metaspace for which you want to view data.

**Space** Select the space for which you want to view data.

#### Fields and Data:

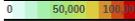
**Last Update** The date and time in which the data in the display was last updated.

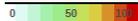
**Member Count** The number of members found for the selected **Domain/Metaspace** combination.

**Labels** Select this check box to display the names of the adapters at the top of each rectangle in the heatmap.

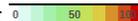
**Log Color Scale** Select this check box to use a logarithmic scale, rather than a linear scale, to map from the selected metric value for a cell to the color for the cell. **Log Scale** provides another way to distribute and differentiate values that you might not be able to see on a linear scale due to the dominant nature of large values in a linear scale.

**Color Metric** Choose a metric to view in the display.

**Entries** The total number of entries in the adapters. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasMemberEntriesHigh**. The middle value in the gradient bar indicates the middle value of the range.

**Gets/sec** The number of gets per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasMemberGetRateHigh**. The middle value in the gradient bar indicates the middle value of the range.

**Puts/sec** The number of puts per second. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasMemberPutRateHigh**. The middle value in the gradient bar indicates the middle value of the range.

**Takes/sec** The number of takes per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasMemberTakeRateHigh**. The middle value in the gradient bar indicates the middle value of the range.

**Expires/sec** The number of expires per second. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasMemberExpireRateHigh**. The middle value in the gradient bar indicates the middle value of the range.

**Evicts/sec** The number of evictions per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasMemberEvictsRateHigh**. The middle value in the gradient bar indicates the middle value of the range.

## Spaces by Member Table

The table in this display provides a view of all of your spaces (by member name) and the their associated metric data including domain, metaspace, space, alert severity, alert count, and the current value of each gathered metric. You can click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by clicking a row to view details for the selected member in the "Member by Space Summary" display.

Domain	Metaspace	MemberName	Space	Alert Severity	Alert Count	Di
Production	retailV212	as-agent-1	customers		0	SEI
Production	retailV212	as-agent-1	inventory		0	SEI
Production	retailV212	as-agent-1	stores		0	SEI

### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the selected domain. Refer to TIBCO ActiveSpaces documentation for more information regarding these fields.

### Filter By:

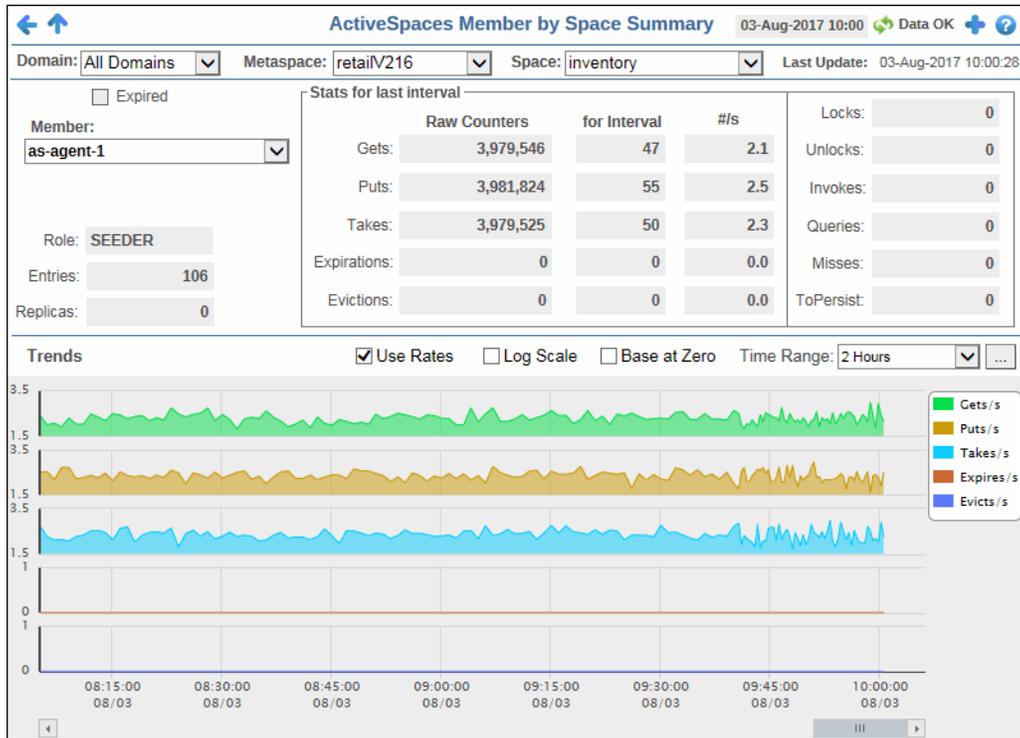
**Domain** Select the domain for which you want to view data.

<b>Metaspace</b>	Select the metaspace for which you want to view data.
<b>Member</b>	Select the space for which you want to view data.
<b>Member Count</b>	The resulting total number of members found in the filtered query, which are listed in the <b>Spaces by Members</b> table.
<b>Spaces by Member Table</b>	
<b>Domain</b>	The name of the domain.
<b>Metaspace</b>	The name of the metaspace.
<b>Member Name</b>	The name of the member.
<b>Space</b>	The name of the space.
<b>Alert Severity</b>	The current alert severity.  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
<b>Alert Count</b>	The total number of alerts for the host.
<b>Distribution Role</b>	The member's role within the space.
<b>Entries</b>	The number of entries.*
<b>% Capacity</b>	The percentage of available entries used for the space.
<b>Replicas</b>	The number of replicas.*
<b>Gets</b>	The total number of "get" operations performed on the user-spaces defined on the metaspace.*
<b>Gets/interval</b>	The number of "get" operations performed on the user-spaces defined for the metaspace during the current polling interval.*
<b>Gets/sec</b>	The rate of "get" operations (per second) performed on the user-spaces defined for the metaspace.*
<b>Puts</b>	The total number of "put" operations performed on the user-spaces defined on the metaspace.*
<b>Puts/interval</b>	The number of "put" operations performed on the user-spaces defined for the metaspace during the current polling interval.*
<b>Puts/sec</b>	The rate of "put" operations (per second) performed on the user-spaces defined for the metaspace.*
<b>Takes</b>	The total number of "take" operations performed on the user-spaces defined on the metaspace.*
<b>Takes/interval</b>	The number of "take" operations performed on the user-spaces defined for the metaspace during the current polling interval.*
<b>Takes/sec</b>	The rate of "take" operations (per second) performed on the user-spaces defined for the metaspace.*
<b>Expires</b>	The total number of entries in the user-spaces defined on the metaspace that have expired.*
<b>Expires/interval</b>	The number of entries performed in the user-spaces defined for the metaspace that expired during the current polling interval.*

<b>Expires/sec</b>	The rate of entries in the user-spaces defined for the metaspace that expired (per second).*
<b>Evicts</b>	The total number of entries in the user-spaces defined on the metaspace that have been evicted.*
<b>Evicts/interval</b>	The number of entries performed in the user-spaces defined for the metaspace that were evicted during the current polling interval.*
<b>Evicts/sec</b>	The rate of entries in the user-spaces defined for the metaspace that were evicted (per second).*
<b>Locks</b>	The total number of locks in the user-spaces defined for the metaspace.*
<b>Unlocks</b>	The total number of unlocks in the user-spaces defined for the metaspace.*
<b>Invokes</b>	The remote invocation count.*
<b>Queries</b>	The total number of queries in the user-spaces defined for the metaspace.*
<b>Misses</b>	The total number of misses in the user-spaces defined for the metaspace.*
<b>ToPersist</b>	The ToPersist count, which indicates how many tuples are required to be persisted to the database if the write-behind feature is configured.*
<b>ClientAvgGetMicros</b>	The client's average "get" latency in microseconds.*
<b>ClientAvgPutMicros</b>	The client's average "put" latency in microseconds.*
<b>ClientAvgTakeMicros</b>	The client's average "take" latency in microseconds.*
<b>ClientMaxGetMicros</b>	The client's highest "get" latency in microseconds.*
<b>ClientMaxPutMicros</b>	The client's highest "put" latency in microseconds.*
<b>ClientMaxTakeMicros</b>	The client's highest "take" latency in microseconds.*
<b>ClientTotalGetMillis</b>	The client's cumulative total "get" latency in milliseconds for the current polling period.*
<b>ClientTotalPutMillis</b>	The client's cumulative total "put" latency in milliseconds for the current polling period.*
<b>ClientTotalTakeMillis</b>	The client's cumulative total "take" latency in milliseconds for the current polling period.*
<b>ClientTotalMissMillis</b>	The client's cumulative total "miss" latency in milliseconds for the current polling period.*
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > (Project Name) > <b>Solution Package Configuration</b> > <b>TIBCO Active Spaces</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Timestamp</b>	The date and time the row data was last updated.

## Member by Space Summary

This display provides a view of the current and historical metrics for a single member in a particular space. The trend graph in the bottom half of the display traces the current and historical total number of or rate data for gets, puts, takes, expires, and evictions.



**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the selected domain. Refer to TIBCO ActiveSpaces documentation for more information regarding these fields.

**Filter By:**

The display might include these filtering options:

- Domain** Select the domain for which you want to show data in the display.
- Metaspace** Select the metaspace for which you want to show data in the display.
- Space** Select the space for which you want to show data in the display.

**Fields and Data:**

- Last Update** The date and time in which the data in the display was last updated.

<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > (Project Name) > <b>Solution Package Configuration</b> > <b>TIBCO Active Spaces</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Member</b>	Select the member for which you want to see data.*
<b>Role</b>	The member's role within the space.*
<b>Entries</b>	The number of entries.*
<b>Replicas</b>	The number of replicas.*
<b>Stats for last interval</b>	
<b>Gets</b>	<b>Raw Counters</b> -- The total number of gets for the space. <b>for interval</b> -- The number of gets for the current interval. <b>#/s</b> -- The number of gets received per second.
<b>Puts</b>	<b>Raw Counters</b> -- The total number of puts for the space. <b>for interval</b> -- The number of puts for the current interval. <b>#/s</b> -- The number of puts received per second.
<b>Takes</b>	<b>Raw Counters</b> -- The total number of takes for the space. <b>for interval</b> -- The number of takes for the current interval. <b>#/s</b> -- The number of takes received per second.
<b>Expirations</b>	<b>Raw Counters</b> -- The total number of expirations for the space. <b>for interval</b> -- The number of expirations for the current interval. <b>#/s</b> -- The number of expirations received per second.
<b>Evictions</b>	<b>Raw Counters</b> -- The total number of evictions for the space. <b>for interval</b> -- The number of evictions for the current interval. <b>#/s</b> -- The number of evictions received per second.
<b>Locks</b>	The total number of locks in the user-spaces defined for the space.*
<b>Unlocks</b>	The total number of unlocks in the user-spaces defined for the space.*
<b>Invokes</b>	The remote invocation count.*
<b>Queries</b>	The total number of queries in the user-spaces defined for the space.*
<b>Misses</b>	The total number of misses in the user-spaces defined for the space.*
<b>ToPersist</b>	The ToPersist count, which indicates how many tuples are required to be persisted to the database if the write-behind feature is configured.*

**Trends**

Traces the following:

**Gets(/s)** -- traces the total number of gets, or the number of gets per second with **Use Rates** selected.

**Puts(/s)**-- traces the total number of puts, or the number of puts per second with **Use Rates** selected.

**Takes(/s)** -- traces the total number of takes, or the number of takes per second with **Use Rates** selected.

**Expires(/s)** -- traces the total number of expires, or the number of expires per second with **Use Rates** selected.

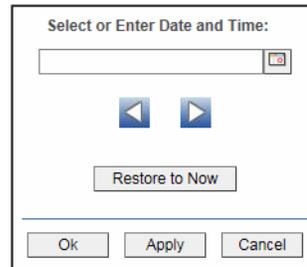
**Evicts(/s)** -- traces the total number of evicts, or the number of evicts per second with **Use Rates** selected.

**Use Rates** Select this check box to trace the rates (**Gets/s, Puts/s, Takes/s, Expires/s, Evicts/s**) instead of the total numbers (**Gets, Puts, Takes, Expires, Evicts**).

**Log Scale** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.

**Base at Zero** Select to use zero (**0**) as the Y axis minimum for all graph traces.

**Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

---

## TIBCO ActiveSpaces (2.x) - HTML

The HTML version features an overview display, "[TIBCO ActiveSpaces 2 Overview](#)" (pictured below), and the following Views which can be found under **Components** tab > **Middleware** > **TIBCO ActiveSpaces 2**:

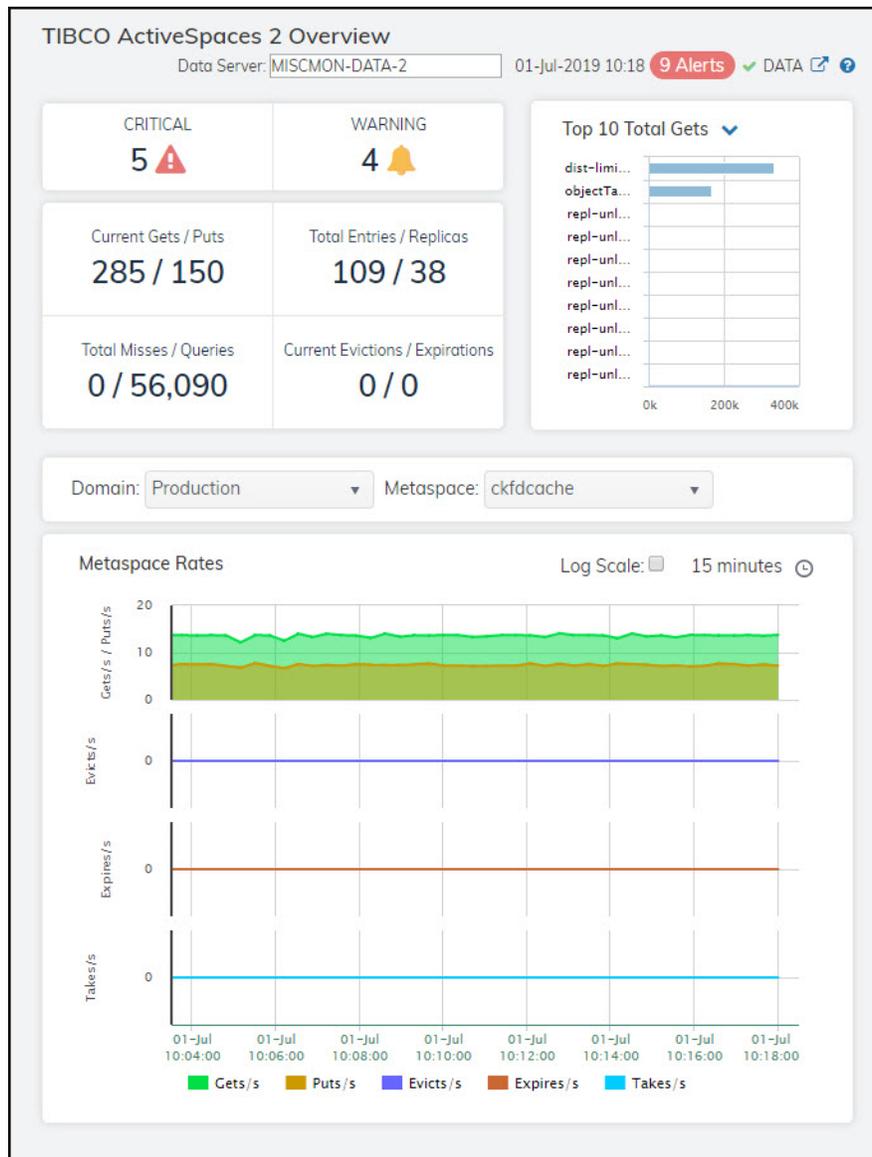
### TIBCO ActiveSpaces 2 Overview

The **TIBCO ActiveSpaces 2 Overview** is the top-level display for the TIBCO ActiveSpaces Monitor, which provides a good starting point for immediately getting the status of all your metaspaces, spaces, and members on your Data Server. You can select the RTView DataServer for which you want to see data and easily view the current data for that DataServer including:

- The total number of active alerts for the selected DataServer, including the total number of critical and warning alerts.
- The current number of current Gets/Puts.
- The number of entries and replicas.
- The number of misses and queries.
- The current number of evictions and expirations.
- A visual list of the top 10 metaspaces containing the total gets/puts/entries/replicas/ expirations on your connected DataServer.

You can hover over each region in the upper half of the Overview to see more detail. You can also drill down to see even more detail by clicking on each respective region in the Overview. For example, clicking on the alerts in the **CRITICAL** and **WARNING** alerts region opens the **Alerts Table by Components** display.

The bottom half of the display provide a **MetaSpace Rates** trend graphs for a selected metaspace. You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.



The following TIBCO ActiveSpaces 2 Views can be found under **Components** tab > **Middleware** > **TIBCO ActiveSpaces**:

- **"MetaSpaces - HTML"**: The displays in this View allow you to view the current and historical metrics for all MetaSpaces in a heatmap, tabular, or summary format.
- **"Spaces - HTML"**: The displays in this View allow you to view the current and historical metrics for all spaces in a heatmap, tabular, or summary format.
- **"Members - HTML"**: The displays in this View allow you to view the current and historical metrics for all members in a particular metaspace, view data for members within a particular space, and view data for all spaces for a particular member.

## MetaSpaces - HTML

These displays present performance metrics and alert status for your TIBCO ActiveSpaces 2 MetaSpaces. Clicking **MetaSpaces** from the left/navigation menu opens the ["TIBCO ActiveSpaces 2 Metaspaces Table - HTML"](#) display, which provides a tabular view of your MetaSpaces and their associated metrics. The option available under **MetaSpaces** is:

- **Single MetaSpace:** Opens the ["TIBCO ActiveSpaces 2 MetaSpace Summary - HTML"](#) display, which shows metrics and trend data for a particular MetaSpace.

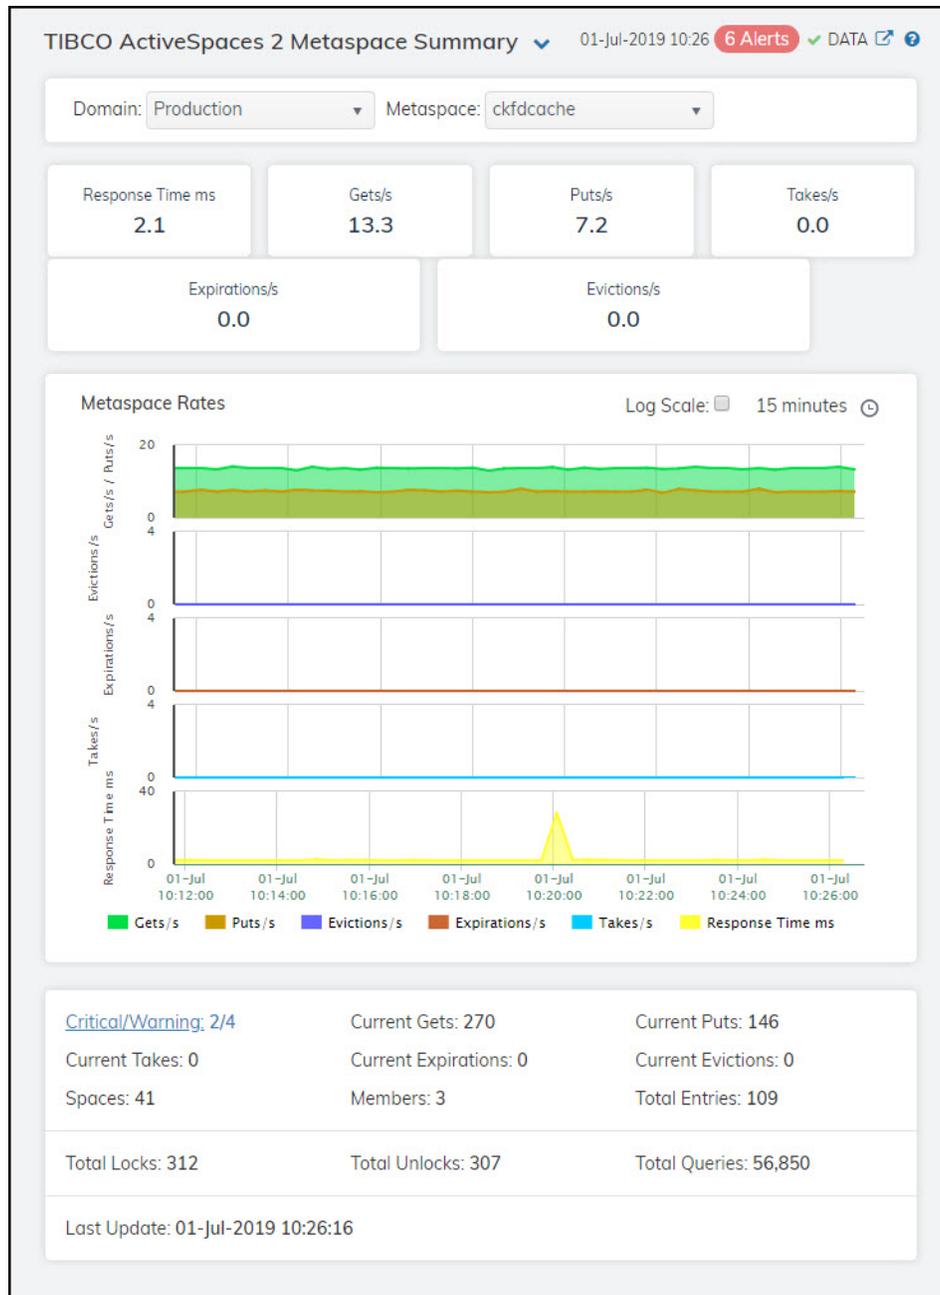
## TIBCO ActiveSpaces 2 Metaspaces Table - HTML

The table in this display provides a view of all of your metaspaces and their associated metric data including domain, members, spaces, alert level, alert count, and the current value of each gathered metric. You can click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by double-clicking a row to view details for the selected metaspace in the ["TIBCO ActiveSpaces 2 MetaSpace Summary - HTML"](#) display.

Domain	Metaspace	Alert Level	Alert Count	Spaces	Members
Production	ckfdcache		6	41	3

## TIBCO ActiveSpaces 2 MetaSpace Summary - HTML

Clicking **Single MetaSpace** in the left/navigation menu opens the **TIBCO ActiveSpaces 2 MetaSpace Summary** display, which allows you to view the current and historical metrics for a single metaspace. Clicking on the information boxes at the top of the display takes you to the ["TIBCO ActiveSpaces 2 Metaspaces Table - HTML"](#) display, where you can view additional metaspace data. The **MetaSpace Rates** trend graph region in the bottom half of the display traces the current and historical total number of or rate data for gets, puts, takes, expires, and evictions, and also traces the average response time. Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display.



## Spaces - HTML

These displays present performance metrics and alert status for your TIBCO ActiveSpaces spaces. Clicking **Spaces** from the left/navigation menu opens the “[TIBCO ActiveSpaces 2 Spaces Table - HTML](#)” display, which shows all available utilization metrics for all spaces within a particular metaspace. The options available under **Spaces** are:

- **All Spaces Heatmap:** Opens the “[TIBCO ActiveSpaces 2 Spaces Heatmap - HTML](#)”, which shows a heatmap view of all spaces contained within a particular metaspace.
- **Space Summary:** Opens the “[TIBCO ActiveSpaces 2 Space Summary - HTML](#)” display, which allows you to view metrics and trend data for a particular space.

## TIBCO ActiveSpaces 2 Spaces Table - HTML

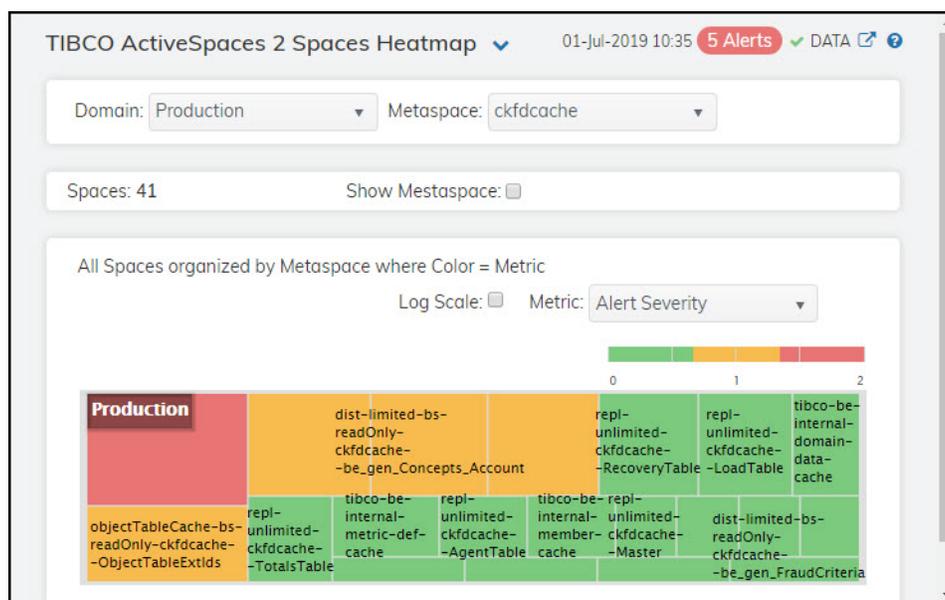
The table in this display provides a view of all of your spaces and their associated metric data including domain, metaspace, space, alert level, alert count, and the current value of each gathered metric. You can click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by double-clicking a row to view details for the selected space in the “[TIBCO ActiveSpaces 2 Space Summary - HTML](#)” display.

Domain	Metaspace	Space	Alert Level	Alert Count	Sp St
Production	ckfdcache	repl-unlimited-ckfdcache--Master	✓	0	READY
Production	ckfdcache	repl-unlimited-ckfdcache--Metadat	✓	0	READY
Production	ckfdcache	repl-unlimited-ckfdcache--LoadTab	✓	0	READY
Production	ckfdcache	dist-limited-nobs-ckfdcache--com_t	✓	0	READY
Production	ckfdcache	dist-limited-hc-readOnly-ckfdcache	✓	0	READY

## TIBCO ActiveSpaces 2 Spaces Heatmap - HTML

Clicking **All Spaces Heatmap** in the left/navigation menu opens the **TIBCO ActiveSpaces 2 Spaces Heatmap**, which provides an easy-to-view interface that allows you to quickly identify the current status of each of your spaces for each available metric. You can view the spaces in the heatmap based on the following metrics: current alert severity, entries, gets per second, puts per second, takes per second, expires per second, and evicts per second. By default, this display shows the heatmap based on the **Alert Severity** metric.

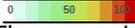
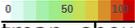
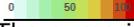
The heatmap is organized so that each rectangle represents a space. The rectangle color indicates the most critical alert state. Click on a node to drill-down to the "[TIBCO ActiveSpaces 2 Space Summary - HTML](#)" display and view metrics for a particular space. Toggle between the commonly accessed displays by clicking the drop down list on the display title. Mouse-over rectangles to view more details about space performance and status.



### Available Metrics

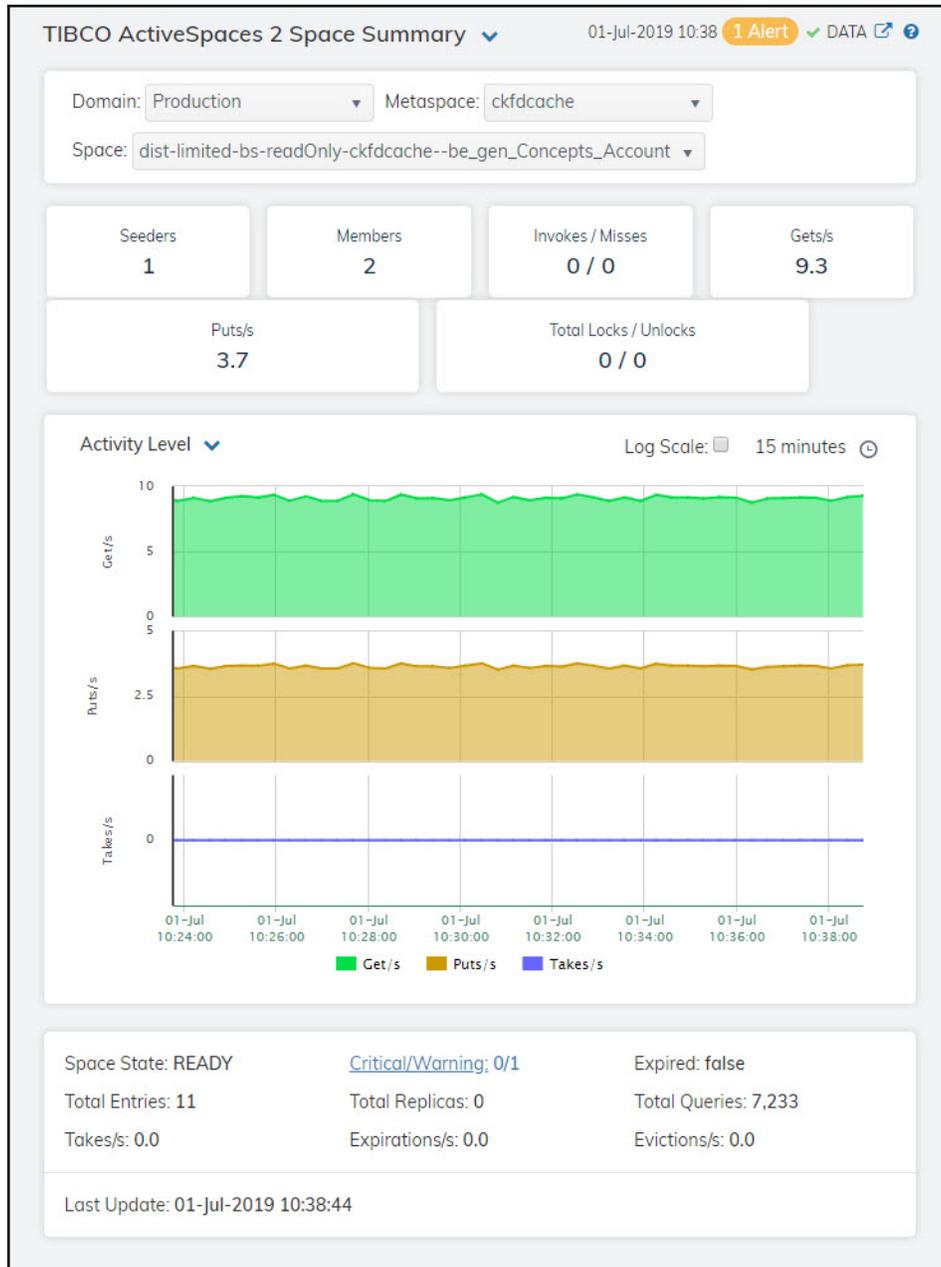
Select the metric driving the heatmap display. The default is **Alert Severity**. Each **Metric** has a color gradient bar that maps values to colors. The heatmap is organized by spaces, where each rectangle represents a space. Mouse-over any rectangle to display the current values of the metrics for the space. Click on a rectangle to drill-down to the associated "[TIBCO ActiveSpaces 2 Space Summary - HTML](#)" display for a detailed view of metrics for that particular space.

- Alert Severity** The current alert severity. Values range from **0** - **2**, as indicated in the color gradient bar, where **2** is the highest Alert Severity:
- Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
  - Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
  - Green indicates that no metrics have exceeded their alert thresholds.
- Alert Count** The total number of critical and warning alerts in the heatmap rectangle. The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.

<b>Entries</b>	The total number of entries in the space. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of <b>TasSpaceEntriesHigh</b> . The middle value in the gradient bar indicates the middle value of the range.
<b>Gets/sec</b>	The number of gets per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of <b>TasSpaceGetRateHigh</b> . The middle value in the gradient bar indicates the middle value of the range.
<b>Puts/sec</b>	The number of message sent per second. The color gradient bar  shows the range of the value/color mapping. ated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of <b>TasSpacePutRateHigh</b> . The middle value in the gradient bar indicates the middle value of the range.
<b>Takes/sec</b>	The number of takes per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of <b>TasSpaceTakeRateHigh</b> . The middle value in the gradient bar indicates the middle value of the range.
<b>Expires/sec</b>	The number of expires per second. The color gradient bar  shows the range of the value/color mapping. ated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of <b>TasSpaceExpireRateHigh</b> . The middle value in the gradient bar indicates the middle value of the range.
<b>Evicts/sec</b>	The number of evictions per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of <b>TasSpaceEvictsRateHigh</b> . The middle value in the gradient bar indicates the middle value of the range.

## TIBCO ActiveSpaces 2 Space Summary - HTML

Clicking **Single Space** in the left/navigation menu opens the **TIBCO ActiveSpaces 2 Space Summary** display, which provides a view of the current and historical metrics for a single space. Clicking on the information boxes at the top of the display takes you to the "[TIBCO ActiveSpaces 2 Spaces Table - HTML](#)", where you can view additional cluster/event data. There are two options in the trend graph region: **Activity Level** and **Exceptions**. In the **Activity Level** trend graph region, you can view "gets" rate, "puts" rate, and "takes" rate over a selected time range. In the **Exceptions** trend graph region, you can view "expires" rate and "evicts" rate over a selected time range. Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display.



## Members - HTML

These displays present performance metrics and alert status for your TIBCO ActiveSpaces 2 spaces. Clicking **Members** from the left/navigation menu opens the “[TIBCO ActiveSpaces 2 All Members Table - HTML](#)” display, which shows a tabular view of all members in a particular metaspace. The options available under **Members** are:

- **All Members Heatmap**: Opens the “[TIBCO ActiveSpaces 2 Members Heatmap - HTML](#)”, which shows a heatmap view of all members in a particular metaspace.
- **Single Member**: Opens the “[TIBCO ActiveSpaces 2 Member Summary - HTML](#)” display, which allows you to view current and trending data for a single member for a particular metaspace.

### TIBCO ActiveSpaces 2 All Members Table - HTML

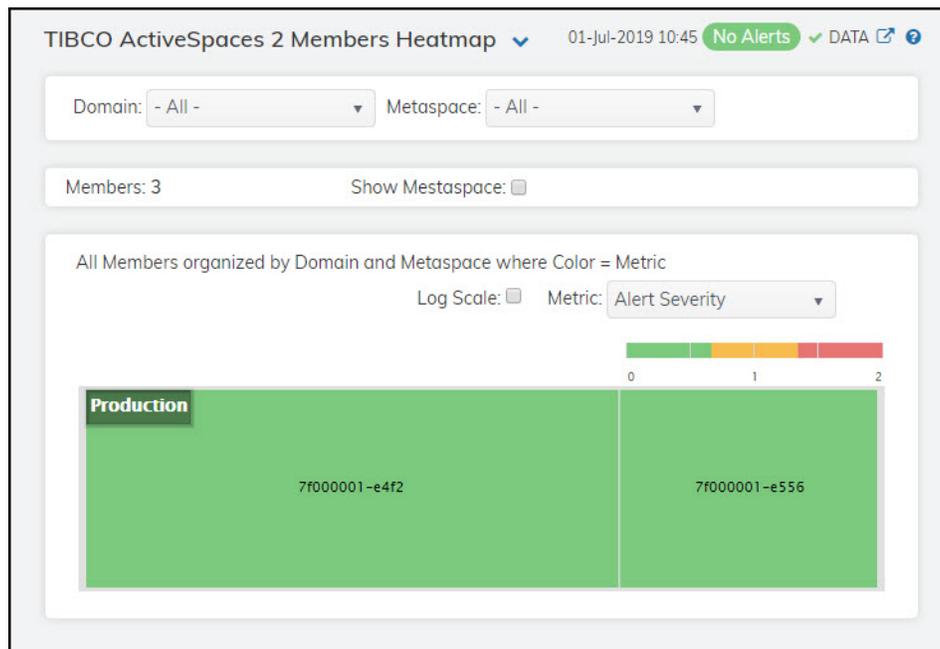
The table in this display provides a view of all of the members in a particular metaspace and their associated metric data including domain, metaspace, alert severity, alert count, and the current value of each gathered metric. You can click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by double-clicking a row to view details for the selected member in the “[TIBCO ActiveSpaces 2 Member Summary - HTML](#)” display.

Domain	Metaspace	Member Name	Alert Level	Alert Count	Management Role
Production	ckfdcache	7f000001-e48e	✓	0	MEMBER
Production	ckfdcache	7f000001-e4f2	✓	0	MEMBER
Production	ckfdcache	7f000001-e556	✓	0	MANAGER

## TIBCO ActiveSpaces 2 Members Heatmap - HTML

Clicking **All Members Heatmap** in the left/navigation menu opens the **TIBCO ActiveSpaces 2 Members Heatmap**, which provides an easy-to-view interface that allows you to quickly identify the current status of each of your members for each available metric. Use the **Metric** drop-down menu to view the heatmap using a different metric.

The heatmap is organized so that each rectangle represents a member. The rectangle color indicates the most critical alert state. Click on a node to drill-down to the ["TIBCO ActiveSpaces 2 Member Summary - HTML"](#) display and view metrics for a particular member. Toggle between the commonly accessed displays by clicking the drop down list on the display title. You can use the **Show Metaspaces** check-box to include or exclude labels in the heatmap, and you can mouse over a rectangle to see additional metrics for a particular member.



### Available Metrics

Select the metric driving the heatmap display. The default is **Alert Severity**. Each **Metric** has a color gradient bar that maps values to colors. The heatmap is organized by members, where each rectangle represents a member. Mouse-over any rectangle to display the current values of the metrics for the member. Click on a rectangle to drill-down to the associated ["TIBCO ActiveSpaces 2 Member Summary - HTML"](#) display for a detailed view of metrics for that particular member.

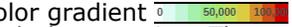
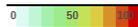
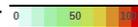
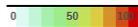
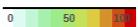
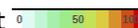
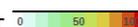
#### Alert Severity

The current alert severity. Values range from **0** - **2**, as indicated in the color gradient bar, where **2** is the highest Alert Severity:

- Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
- Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
- Green indicates that no metrics have exceeded their alert thresholds.

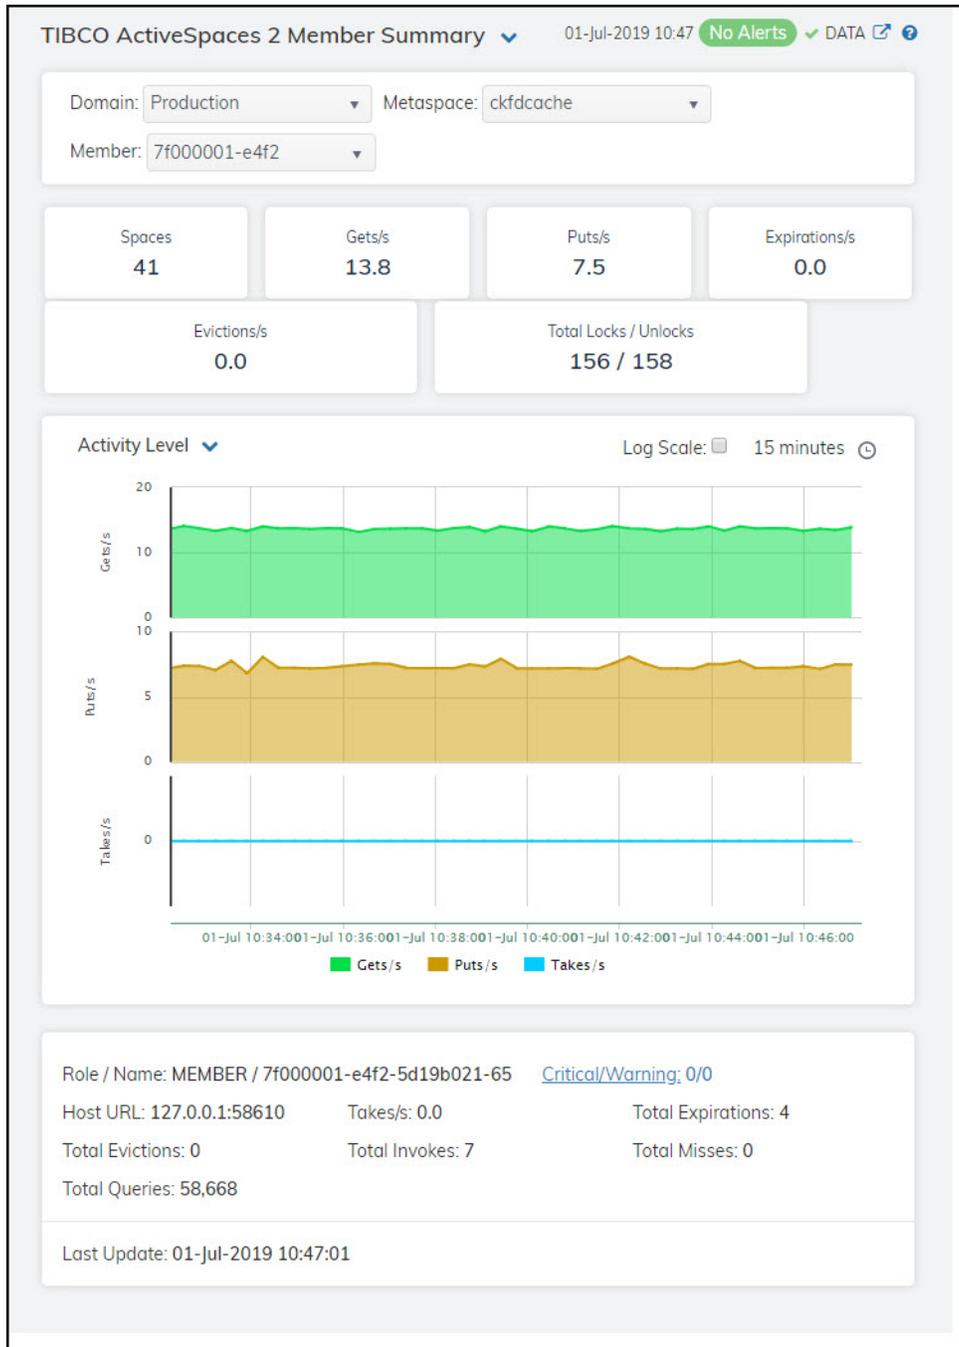
#### Alert Count

The total number of critical and warning alerts in the heatmap rectangle. The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.

- Entries** The total number of entries for the member. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasMemberEntriesHigh**. The middle value in the gradient bar indicates the middle value of the range.
- Gets/sec** The number of gets per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasMemberGetRateHigh**. The middle value in the gradient bar indicates the middle value of the range.
- Puts/sec** The number of puts per second. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasMemberPutRateHigh**. The middle value in the gradient bar indicates the middle value of the range.
- Takes/sec** The number of takes per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasMemberTakeRateHigh**. The middle value in the gradient bar indicates the middle value of the range.
- Expires/sec** The number of expires per second. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasMemberExpireRateHigh**. The middle value in the gradient bar indicates the middle value of the range.
- Evicts/sec** The number of evictions per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasMemberEvictsRateHigh**. The middle value in the gradient bar indicates the middle value of the range.
- CPU %** The percentage of CPU used. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasMemberCpuHigh**. The middle value in the gradient bar indicates the middle value of the range.
- Memory %** The percentage of memory used. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasMemberMemoryUsedHigh**. The middle value in the gradient bar indicates the middle value of the range.
- JVM Memory %** The percentage of JVM memory used. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TasMemberJvmMemoryUsedHigh**. The middle value in the gradient bar indicates the middle value of the range.

## TIBCO ActiveSpaces 2 Member Summary - HTML

Clicking **Single Member** in the left/navigation menu opens the **TIBCO ActiveSpaces 2 Member Summary** display, which allows you to view of the current and historical metrics for a single member. Clicking on the information boxes at the top of the display takes you to the ["TIBCO ActiveSpaces 2 All Members Table - HTML"](#) display, where you can view additional member data. There are two options in the trend graph region: **Activity Level** and **Exceptions**. In the **Activity Level** trend graph region, you can view "gets" rate, "puts" rate, and "takes" rate over a selected time range. In the **Exceptions** trend graph region, you can view "expirations" rate and "evictions" rate over a selected time range. Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display.



---

## TIBCO Adapters

The following TIBCO Adapters Views can be found under **Components** tab > **Middleware** > **TIBCO Adapters**. The displays within the Views will be populated with data once the Solution Package for TIBCO Adapters is configured in the RTView DataServer for TIBCO and the RTView DataServer for TIBCO is connected to RTViewCentral.

- **"All Adapters View"**: The displays in this View allow you to view the current and historical metrics for all adapters in a heatmap or tabular format.
- **"Single Adapter View"**: The displays in this View allow you to view the current and historical metrics for a single adapter in a tabular format.

### All Adapters View

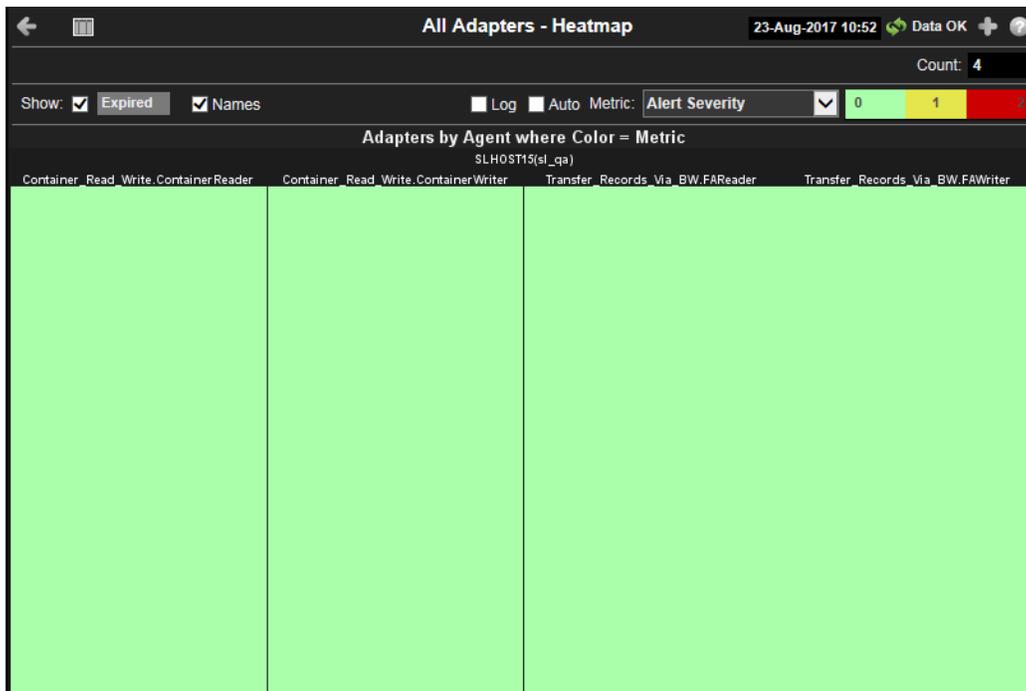
These displays provide detailed data for all adapters. Displays in this View are:

- **"All Adapters Heatmap"**: A heatmap view of all adapters in a heatmap format and their associated metrics.
- **"All Adapters Table"**: A tabular view of your adapters and their associated metrics.

### All Adapters Heatmap

This heatmap display provides an easy-to-view interface that allows you to quickly identify the current status of each of your adapters for each available metric. You can view the adapters in the heatmap based on the following metrics: the current alert severity, the current alert count, the delta messages received, the messages received rate, the messages sent rate, the delta messages sent, and the increase in errors from the previous polling period. By default, this display shows the heatmap based on the **Alert Severity** metric.

You can use the **Names** check-box  to include or exclude labels in the heatmap, and you can mouse over a rectangle to see additional metrics for an adapter. Clicking one of the rectangles in the heatmap opens the **"Adapter Summary"** display, which allows you to see additional details for the selected adapter



#### Title Bar (possible features are):

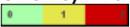
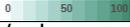
- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

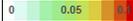
Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

#### Fields and Data:

- Count** The number of adapters included in the display. This number can change if you toggle the **Expired** check box on and off.
- Show: Expired** Select this check box to display those adapters whose data has not been updated recently (expired).
- Show: Names** Select this check box to display the names of the adapters at the top of each rectangle in the heatmap.
- Log** Select this check box to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Auto** Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value.  
**Note:** Some metrics auto-scale automatically, even when **Auto** is not selected.
- Metric** Choose a metric to view in the display.

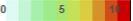
<b>Alert Severity</b>	<p>The current alert severity. Values range from <b>0</b> - <b>2</b>, as indicated in the color gradient  bar, where <b>2</b> is the highest Alert Severity:</p> <ul style="list-style-type: none"> <li> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Alert Count</b>	<p>The total number of critical and warning unacknowledged alerts in the adapters. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.</p>
<b>Delta Msgs Rcvd</b>	<p>The increase in the number of messages received (per second) from the previous polling period to the current polling period. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum number of messages received. The middle value in the gradient bar indicates the middle value of the range.</p> <p>The <b>Auto</b> check box does not impact this metric.</p>
<b>Rate Msgs Rcvd</b>	<p>The number of messages received per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of <b>TadAdapterMsgsRcvdRateHigh</b>. The middle value in the gradient bar indicates the middle value of the range.</p> <p>When <b>Auto</b> is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>
<b>Delta Msgs Sent</b>	<p>The increase in the number of messages sent (per second) from the previous polling period to the current polling period. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum number of messages sent. The middle value in the gradient bar indicates the middle value of the range.</p> <p>The <b>Auto</b> check box does not impact this metric.</p>

**Rate Msgs Sent**

The number of message sent per second. The color gradient bar  shows the range of the value/color mapping. The current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TadAdapterMsgsSentRateHigh**. The middle value in the gradient bar indicates the middle value of the range.

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

**Delta Errors**

The increase in the number of errors from the previous polling period to the current polling period. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TadAdapterDeltaErrorsHigh**. The middle value in the gradient bar indicates the middle value of the range.

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

**All Adapters Table**

The table in this display provides a view of all of your adapters and their associated metric data including agent, application instance, alert severity, alert count, and the current value of each gathered metric. You can click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by clicking a row to view details for the selected adapter in the ["Adapter Summary"](#) display

Agent	Application Instance	Alert Level	Alert Count	Messages Received	Delta Msgs Received	Rate
SLHOST15(sl_qa)	Transfer_Records_Via_BW.FAWriter	Green		0	0	
SLHOST15(sl_qa)	Container_Read_Write.ContainerReader	Green		0	0	
SLHOST15(sl_qa)	Transfer_Records_Via_BW.FAReader	Green		0	0	
SLHOST15(sl_qa)	Container_Read_Write.ContainerWriter	Green		0	0	

**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- 6,047** The number of items currently in the display.
- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the selected adapter. Refer to TIBCO Adapter documentation for more information regarding these fields.

**Fields and Data:**

**Show: Expired** Select this check box to display adapters that have expired data in the table.

**Count** The total number of adapters listed in the **All Adapters Table**.

**All Adapters Table:**

**Agent** The name of the agent.

**Application Instance** The name of the application instance.

<b>Alert Level</b>	The current alert severity.  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
<b>Alert Count</b>	The total number of alerts for the host.
<b>Messages Received</b>	The number messages received.*
<b>Delta Messages Received</b>	The increase in the number of messages received (from the previous polling period to the current polling period).*
<b>Rate Messages Received</b>	The number of messages received per second.*
<b>Messages Sent</b>	The total number of messages sent.*
<b>Delta Msgs Sent</b>	The increase in the number of messages sent (from the previous polling period to the current polling period).*
<b>Rate Msgs Sent</b>	The number of messages sent per second.*
<b>New Errors</b>	The number of new errors received since the last polling update.*
<b>Total Errors</b>	The total number of errors.*
<b>Delta Total Errors</b>	The increase in the number total errors (from the previous polling period to the current polling period).*
<b>Rate Total Errors</b>	The number of errors per second.*
<b>Adapter Name</b>	The name of the adapter.*
<b>Last Restart</b>	The date and time the adapter was last restarted.*
<b>Process ID</b>	The process ID of TIBCO Adapter you are running.*
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > ( <b>Project Name</b> ) > <b>Solution Package Configuration</b> > <b>TIBCO Adapters</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Timestamp</b>	The date and time the row data was last updated.

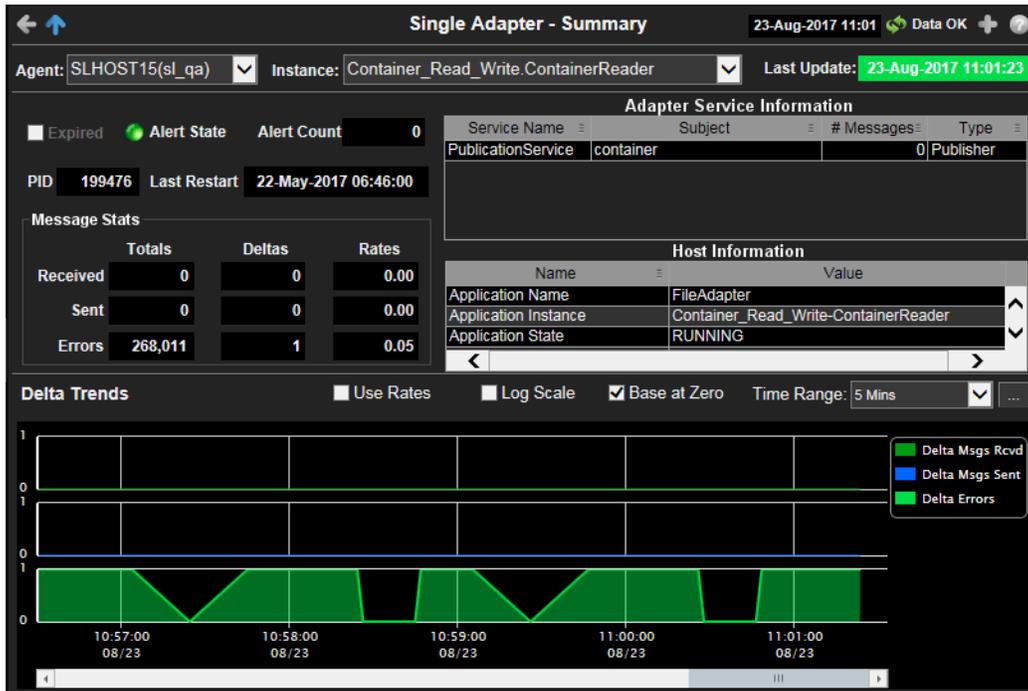
## Single Adapter View

This display allows you to view the current and historical metrics for a single adapter. The available display in this View is:

- **"Adapter Summary"**: This display allows you to view current and trending data for a single adapter for a particular agent.

## Adapter Summary

This display provides a view of the current and historical metrics for a single adapter. You can view message statistics, adapter service information, and host information for a specific instance. The trend graph in the bottom half of the display traces the current and historical delta messages received, delta messages sent, and delta errors.



**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu Table open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by the selected adapter. Refer to TIBCO Adapter documentation for more information regarding these fields.

**Filter By:**

The display might include these filtering options:

- Agent** Select the agent for which you want to show data in the display.
- Instance** Select the instance for which you want to show data in the display.

**Fields and Data:**

<b>Last Update</b>	The date and time in which the data in the display was last updated.
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > ( <b>Project Name</b> ) > <b>Solution Package Configuration</b> > <b>TIBCO Adapters</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Alert State</b>	The current alert severity.  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
<b>Alert Count</b>	The total number of current alerts.
<b>PID</b>	The process ID of the Instance.*
<b>Last Restart</b>	The data and time the instance was last restarted.*

**Message Stats**

<b>Received</b>	<b>Totals</b> -- The total number of messages received.* <b>Deltas</b> -- The increase in the number of messages received since the last polling update.* <b>Rates</b> -- The number of messages received per second.*
<b>Sent</b>	<b>Totals</b> -- The total number of messages sent.* <b>Deltas</b> -- The increase in the number of messages sent since the last polling update.* <b>Rates</b> -- The number of messages sent per second.*
<b>Errors</b>	<b>Totals</b> -- The total number of errors that have occurred.* <b>Deltas</b> -- The increase in the number of errors since the last polling update.* <b>Rates</b> -- The number of errors occurring per second.*

**Adapter Service Information**

<b>Service Name</b>	The name of the service.*
<b>Subject</b>	The name of the subject.*
<b># Messages</b>	The current number of messages.*
<b>Type</b>	The type of adapter service.*

**Host Information**

<b>Name</b>	The name of the host.*
<b>Value</b>	The host's value.*

**Trends Graph** Traces the following:

**Delta Msgs Rcvd** -- traces the increase in the number of messages received since the last polling update, or the rate of messages received with **Use Rates** selected.

**Delta Msgs Sent** -- traces the increase in the number of messages sent since the last polling update, or the rate of messages sent with **Use Rates** selected.

**Delta Errors** -- traces the increase in the number of errors since the last polling update, or the rate of errors with **Use Rates** selected.

**Use Rates** Select this toggle to trace the rates (**Msgs Rcvd/sec**, **Msgs Sent/sec**, and **Errors/sec**) instead of the delta numbers (**Delta Msgs Rcvd**, **Delta Msgs Sent**, and **Delta Errors**).

**Log Scale** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.

**Base at Zero** Select to use zero (0) as the Y axis minimum for all graph traces.

**Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .

By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

---

## TIBCO Adapters - HTML

The following Views and their associated displays are available in the Monitor. This section describes the Monitor displays and includes:

- [“TIBCO Adapters Overview - HTML”](#): Describes the **TIBCO Adapters Overview** display.
- [“All Adapters View - HTML”](#): The displays in this View allow you to view the current and historical metrics for all adapters in a heatmap or tabular format.

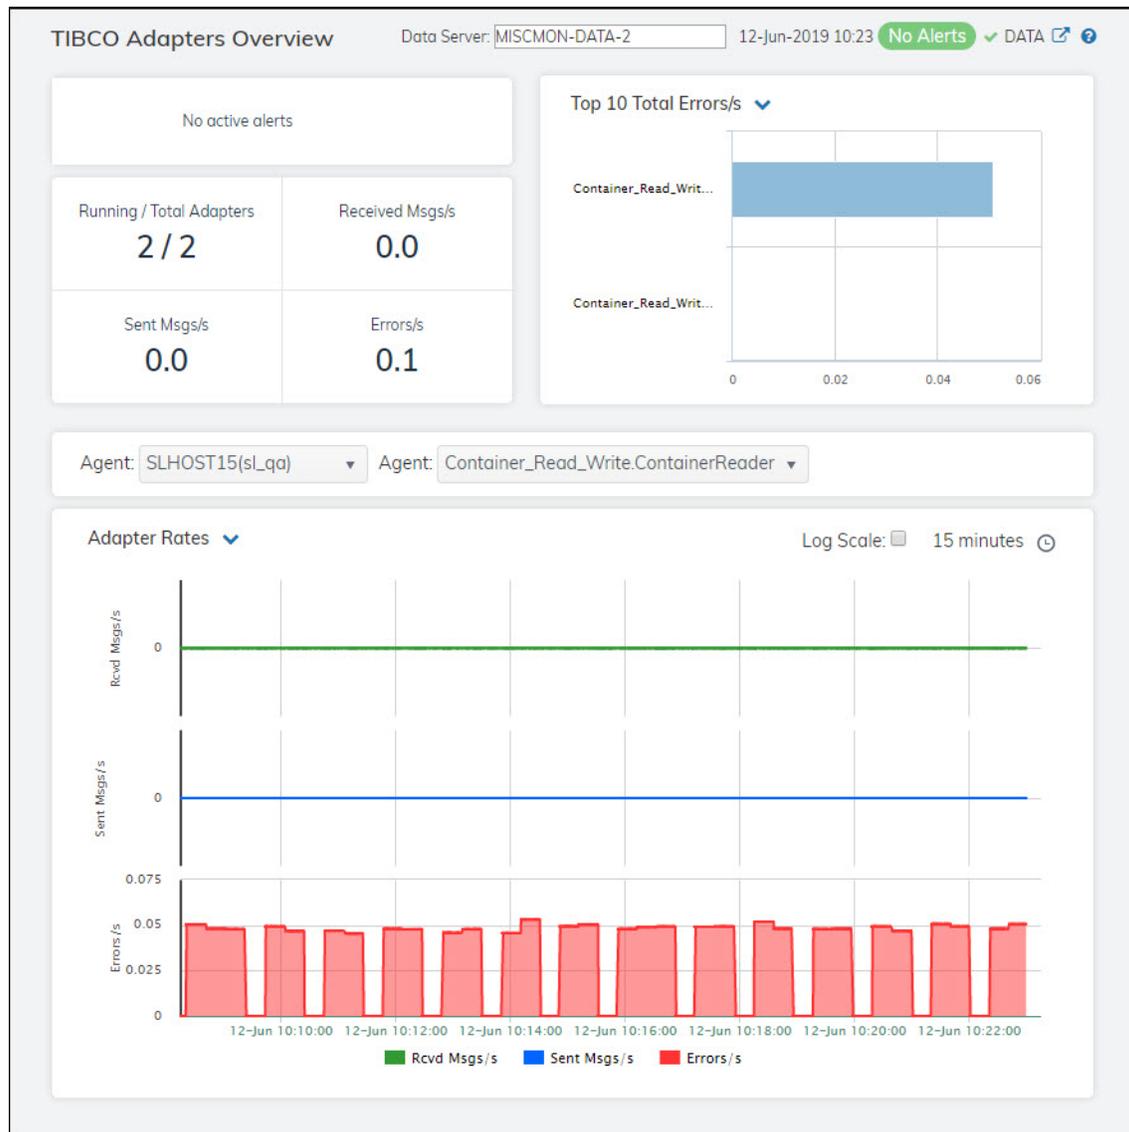
### TIBCO Adapters Overview - HTML

The **TIBCO Adapters Overview** is the top-level display for the TIBCO Adapters Monitor, which provides a good starting point for immediately getting the status of all your adapters on your Data Server. You can select the RTView DataServer for which you want to see data and easily view the current data for that DataServer including:

- The total number of active alerts for the selected DataServer, including the total number of critical and warning alerts.
- The total number of adapters and the number of adapters that are up and running.
- The rate of messages being sent and received.
- The number of errors occurring per second.
- A visual list of the top 10 adapters based on rate of messages being received, rate of messages being sent, and rate of errors occurring on your connected DataServer.

You can hover over each region in the upper half of the Overview to see more detail. You can also drill down to see even more detail by clicking on each respective region in the Overview. For example, clicking on the alerts in the **CRITICAL** and **WARNING** alerts region opens the **Alerts Table by Components** display.

The bottom half of the display provides a host resources trend graph for a selected agent. You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.



## All Adapters View - HTML

These displays provide detailed data for all adapters or for a particular adapter. Clicking **All Adapters** from the left/navigation menu opens the "TIBCO Adapters Table - HTML" display, which shows a tabular view of your adapters and their associated metrics. The options available under **All Adapters** are:

- **All Adapters Heatmap:** Opens the "TIBCO Adapters Heatmap - HTML" display, which provides a view of all adapters in a heatmap format and their associated metrics.
- **Single Adapter:** Opens the "TIBCO Adapter Summary - HTML" display, which provides current and historical data for a single adapter.

## TIBCO Adapters Table - HTML

The table in this display provides a view of all of your adapters and their associated metric data including agent, application instance, alert severity, alert count, and the current value of each gathered metric. You can click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by double-clicking a row to view details for the selected adapter in the ["TIBCO Adapter Summary - HTML"](#) display

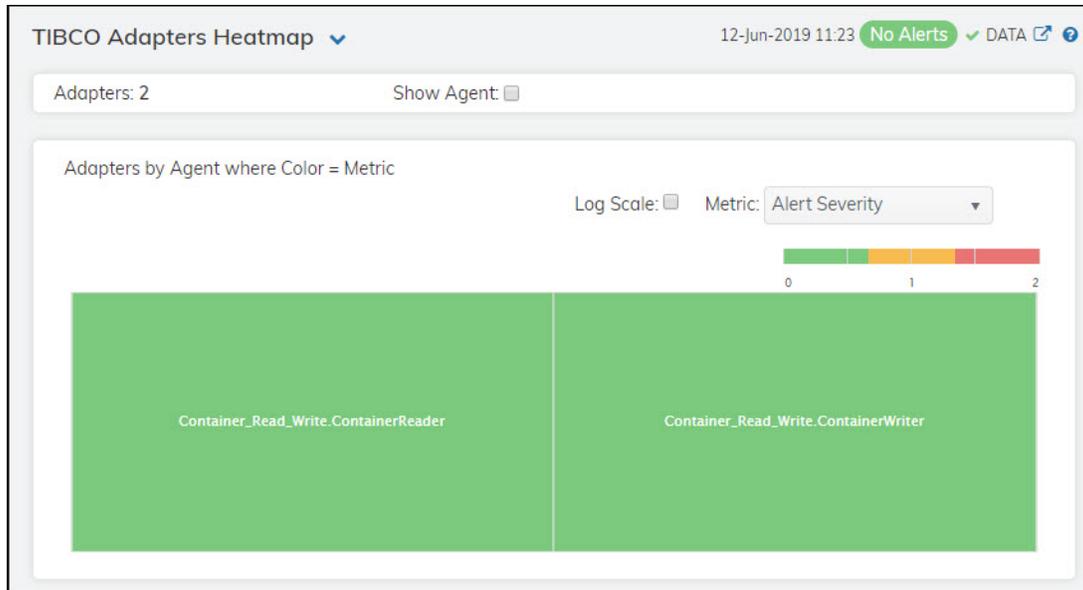
Agent Name	Application Instance	Expired	Alert Level	Alert Count
SLHOST15(sl_qa)	Container_Read_Write.ContainerReader		✓	0
SLHOST15(sl_qa)	Container_Read_Write.ContainerWriter		✓	0

## TIBCO Adapters Heatmap - HTML

Clicking **All Adapters Heatmap** in the left/navigation menu opens the **TIBCO Adapters Heatmap**, which provides an easy-to-view interface that allows you to quickly identify the current status of each of your adapters for each available metric. You can view the adapters in the heatmap based on the following metrics: the current alert severity, the current alert count, the messages received rate, the messages sent rate, the current messages sent (since the last data update), the current messages received (since the last data update), and the increase in errors (since the last data update). By default, this display shows the heatmap based on the **Alert Severity** metric.

Each rectangle in the heatmap represents an adapter. The rectangle color indicates the most critical alert state associated with the adapter. Choose a different metric to display from the **Metric** drop-down menu. Mouse over a rectangle to see additional metrics. By default, this display shows **Alert Severity**.

You can use the **Show Agent** check-box  to include or exclude labels in the heatmap, and you can mouse over a rectangle to see additional metrics for an adapter. Drill-down and investigate an engine by clicking a rectangle in the heatmap to view details in the ["TIBCO Adapter Summary - HTML"](#) display.



### Available Metrics

Select the metric driving the heatmap display. The default is **Alert Severity**. Each **Metric** has a color gradient bar that maps values to colors. The heatmap organizes the servers by agent, where each rectangle represents an adapter. Mouse-over any rectangle to display the current values of the metrics for the adapter. Click on a rectangle to drill-down to the associated ["TIBCO Adapter Summary - HTML"](#) display for a detailed view of metrics for that particular adapter.

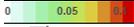
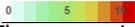
**Alert Severity** The current alert severity. Values range from **0** - **2**, as indicated in the color gradient bar, where **2** is the highest Alert Severity:

- Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
- Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
- Green indicates that no metrics have exceeded their alert thresholds.

**Alert Count** The total number of critical and warning unacknowledged alerts in the adapters. The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.

**Sent Msgs/s** The number of message sent per second. The color gradient bar shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TadAdapterMsgsSentRateHigh**. The middle value in the gradient bar indicates the middle value of the range.

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

<b>Received Msgs/s</b>	<p>The number of messages received per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of <b>TadAdapterMsgsRcvdRateHigh</b>. The middle value in the gradient bar indicates the middle value of the range.</p> <p>When <b>Auto</b> is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>
<b>Current Sent Msgs</b>	<p>The increase in the number of messages sent (per second) from the previous polling period to the current polling period. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum number of messages sent. The middle value in the gradient bar indicates the middle value of the range.</p> <p>The <b>Auto</b> check box does not impact this metric.</p>
<b>Current Received Msgs</b>	<p>The increase in the number of messages received (per second) from the previous polling period to the current polling period. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum number of messages received. The middle value in the gradient bar indicates the middle value of the range.</p> <p>The <b>Auto</b> check box does not impact this metric.</p>
<b>Current Errors</b>	<p>The increase in the number of errors from the previous polling period to the current polling period. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of <b>TadAdapterDeltaErrorsHigh</b>. The middle value in the gradient bar indicates the middle value of the range.</p> <p>When <b>Auto</b> is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>

## TIBCO Adapter Summary - HTML

Clicking **Single Adapter** in the left/navigation menu opens the **TIBCO Adapter Summary** display, which allows you to view current as well as trending data for a single adapter. Clicking on the information boxes at the top of the display takes you to the "[TIBCO Adapters Table - HTML](#)" display, where you can view additional adapter data. You can view message statistics, adapter service information, and host information for a specific instance.

The trend graph has two options: **Adapter Rates** and **Adapter Current Counts**. **Adapter Rates** traces the current and historical rate of messages received, rate of messages sent, and rate of errors over a selected time range. **Adapter Current Counts** traces current and historical messages received, messages sent, and number of errors over a selected time range.

Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display.

TIBCO Adapter Summary ▼
12-Jun-2019 11:37 No Alerts ✓ DATA 🔗 🔔

Agent: SLHOST15(sl\_qa) ▼
Instance: Container\_Read\_Write.ContainerReader ▼

Received Msgs/s  
**0.0**

Sent Msgs/s  
**0.0**

Errors/s  
**0.1**

Current Received Msgs  
**0**

Current Sent Msgs  
**0**

Current Errors  
**1**

#### Adapter Service Information

Service Name	Subject	Type	# Messages
PublicationService	container	Publisher	0

#### Host Information

Name	Value
Application Name	FileAdapter
Application Instance	Container_Read_Write-ContainerReader
Application State	RUNNING
hostName	SLHOST15

#### Adapter Rates ▼

Log Scale:  15 minutes 🕒

■ Rcvd Msgs/s    ■ Sent Msgs/s    ■ Errors/s

Uptime: 9 seconds
[Critical/Warning: 0/0](#)
New Errors: 1

Last Update: 12-Jun-2019 11:37:31

---

## TIBCO BusinessEvents

The following TIBCO BusinessEvents Views can be found under **Components** tab > **Middleware**. The displays within the Views will be populated with data once the Solution Package for TIBCO BusinessEvents is configured in the RTView DataServer for TIBCO and the RTView DataServer for TIBCO is connected to RTViewCentral.

This section includes:

- ["Clusters / Nodes View"](#)
- ["Events / Concepts View"](#)

### Clusters / Nodes View

These displays present performance data for your BusinessEvents system. Displays in this View are:

- ["Clusters"](#)
- ["Cluster Summary"](#)
- ["Cluster Nodes Table"](#)
- ["Cluster Nodes Heatmap"](#)
- ["Inference Node Summary"](#)
- ["Storage Node Summary"](#)

### Clusters

Use this display to check event, concept, and backing store metrics for all of your clusters. Consider keeping this display open to monitor your TIBCO BusinessEvents clusters in general. Each row in the table is a different cluster. Click on a cluster row to view additional cluster details (current and historical) in the ["Cluster Summary"](#) display. The summary display includes trend charts so that you can view key metrics over time.

Sort  the table columns when all the rows cannot fit on the screen. For example, sort the **Alert Status** column so that all nodes with red alerts (●) are listed at the top, or sort the **Expired** column so that all expired nodes are listed at the top.



Cluster Name	Alert Severity	Alert Count	Member Count	Num Events Received	Num Events Sent	Events Received Per Sec	Num Asserted From Channel Per Sec	Num Retracted From Channel Per
ckfdcache		0	2	182,406	0	4.52	4.51	
fdcache		0	2	0	0	0.00	0.00	

**Title Bar** (possible features are):

-   Open the previous and upper display.
-  Open an instance of this display in a new window.
-  Open the online help page for this display.
- open commonly accessed displays.
- The number of items currently in the display.
-  **Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
-  Open the **Alert Views - RTView Alerts Table** display.

**Clusters Table**

Each row in the table is a different cluster, and data in the row columns describe the cluster.

**Note:** Row columns in this table with an asterisk (\*) at the end of the column definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these columns.

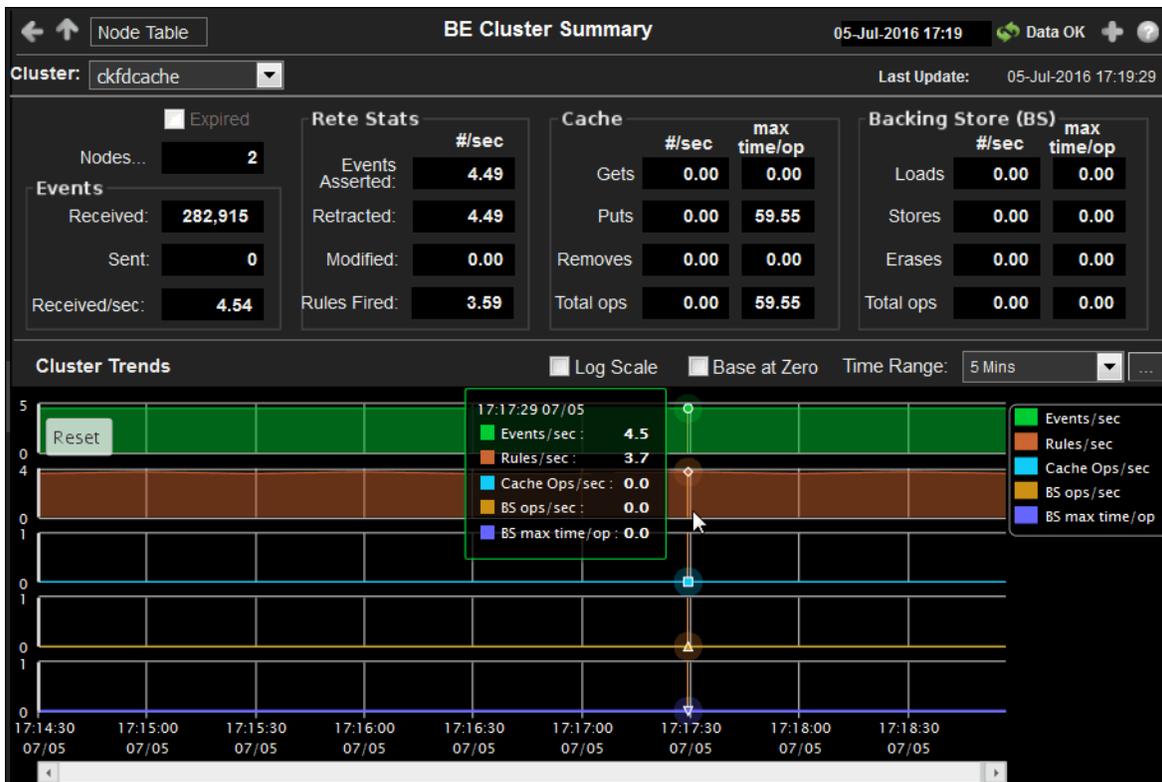
- Count:** The total number of clusters in the table.
- Cluster Name** The name of the TIBCO BusinessEvents cluster.
- Alert Severity** The severity level of open alerts. Values range from **0** to **2**, where **2** is the greatest Severity:
  -  One or more alerts exceeded their ALARM LEVEL threshold.
  -  One or more alerts exceeded their WARNING LEVEL threshold.
  -  No alert thresholds have been exceeded.

<b>Alert Count</b>	The total number of critical and warning alerts.
<b>Member Count</b>	<p>The count of the number of nodes (both cache and inference) that have been collected. For example, for a cluster that has 3 inference nodes and two cache nodes, the Member Count for all 5 rows in the Cluster Table should be 5. If one of the rows shows a member count of one and the others show four, that is a clear indication that a node failed to join the cluster, and the corresponding node should be restarted.</p> <p><b>Note:</b> The actual number of nodes in the cluster will not match the count in this column if one or more of the nodes do not have connection properties configured in the property file that is read by the data server at startup.</p>
<b>Num Events Received</b>	The total number of events received.*
<b>Num Events Sent</b>	The total number of events sent.*
<b>Events Received Per Sec</b>	The rate of events received in the cluster.
<b>Num Asserted From Channel Per Sec</b>	The rate of events asserted into the Rete network via the channel.
<b>Num Retracted From Channel Per Sec</b>	The rate of events retracted/deleted from the Rete network via the channel.
<b>Num Modified From Channel Per Sec</b>	The rate of events modified in the Rete network via the channel.
<b>Num Rules Fired Rate</b>	The rate of rules fired in the cluster.
<b>Concept Max Get Time</b>	The longest time taken for a "get" operation for any node in the cluster since the cluster was started.*
<b>Concept Max Put Time</b>	The longest time taken for a "put" operation for any node in the cluster since the cluster was started.*
<b>Concept Max Remove Time</b>	The longest time taken for a "remove" operation for any node in the cluster since the cluster was started.*
<b>Concept Max Operation Time</b>	The longest time taken for a concept operation (get/put/remove) for any node in the cluster since the cluster was started.*
<b>Concept Gets/sec</b>	The rate of "get" operations in the cluster.
<b>Concept Puts/sec</b>	The rate of "put" operations in the cluster.
<b>Concept Removes/sec</b>	The rate of "remove" operations in the cluster.
<b>Concept Operations/sec</b>	The rate of operations (gets/puts/removes) in the cluster.
<b>Backing Store Max Erase Time</b>	The longest time taken for an "erase" operation in the Backing Store for any node in the cluster.*
<b>Backing Store Max Load Time</b>	The longest time taken for a "load" operation in the Backing Store for any node in the cluster.*
<b>Backing Store Max Store Time</b>	The longest time taken for a "store" operation in the Backing Store for any node in the cluster.*

<b>Backing Store Max Operation Time</b>	The longest time taken to perform an operation (erase/load/store) in the Backing Store for any node in the cluster.*
<b>Backing Store Erases/sec</b>	The rate of "erases" in the Backing Store.
<b>Backing Store Loads/sec</b>	The rate of "loads" into the Backing Store.
<b>Backing Store Stores/sec</b>	The rate of "stores" into the Backing Store.
<b>Backing Store Operations/sec</b>	The rate of operations (erases/loads/stores) in the Backing Store.
<b>Source</b>	The name of the data server from which the data was collected.
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > ( <b>Project Name</b> ) > <b>Solution Package Configuration</b> > <b>TIBCO BusinessEvents</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Timestamp</b>	The date and time, relative to the Data Server, that data was last collected for the engine.

## Cluster Summary

Use this display to view configuration and utilization data for a single cluster. Select a cluster to view Rete statistics, cache metrics, Backing Store data, and trend data for the cluster.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- The number of items currently in the display.

Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

Open the Alert Views - RTView Alerts Table display.

### Filter By:

The display might include these filtering options:

- Cluster** Choose a cluster for which you want to see metrics.
- Last Update** The date and time the data was last updated in the display.

### Fields and Data

This display includes:

**Note:** Fields with an asterisk (\*) at the end of the field definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these fields.

<b>Expired</b>		When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > ( <b>Project Name</b> ) > <b>Solution Package Configuration</b> > <b>TIBCO BusinessEvents</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Nodes</b>		Lists the number of nodes on the cluster.
<b>Events</b>	<b>Received</b>	The number of events received since the last data update.*
	<b>Sent</b>	The number of events sent since the last data update.*
	<b>Received/sec</b>	The rate of events received in the cluster.
<b>Rete Stats</b>	<b>Events Asserted (#/sec)</b>	The rate of events asserted into the Rete network.
	<b>Retracted (#/sec)</b>	The rate of events retracted/deleted from the Rete network.
	<b>Modified (#/sec)</b>	The rate of events modified in the Rete network.
	<b>Rules Fired (#/sec)</b>	The rate of rules fired in the Rete network.
<b>Cache</b>	<b>Gets (#/sec)</b>	The rate of "get" operations in the L1 cache.
	<b>Gets (max time/op)</b>	The longest time taken for a "get" operation for any node in the cluster since the cluster was started.*
	<b>Puts (#/sec)</b>	The rate of "put" operations in the cache.
	<b>Puts (max time/op)</b>	The longest time taken for a "put" operation for any node in the cluster since the cluster was started.*
	<b>Removes (#/sec)</b>	The rate of "removes" in the cache.
	<b>Removes (max time/op)</b>	The longest time taken for a "remove" operation for any node in the cluster since the cluster was started.*
	<b>Total ops (#/sec)</b>	The rate of operations (gets/puts/removes) in the cache.
	<b>Total ops (max time/op)</b>	The longest time taken for an operation (get/put/remove) for any node in the cluster since the cluster was started.*
<b>Backing Store (BS)</b>	<b>Loads (#/sec)</b>	The rate of "load" operations into the backing store.
	<b>Loads (max time/op)</b>	The longest time taken for a "load" operation in the backing store for any node in the cluster.*
	<b>Stores (#/sec)</b>	The rate of "store" operations in the backing store.

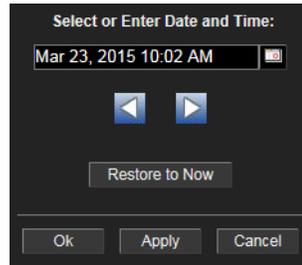
<b>Stores (max time/op)</b>	The longest time taken for a "store" operation in the backing store for any node in the cluster.*
<b>Erases (#/sec)</b>	The rate of "erase" operations in the backing store.
<b>Erases (max time/op)</b>	The longest time taken for an "erase" operation in the backing store for any node in the cluster.*
<b>Total ops (#/sec)</b>	The rate of operations (loads/stores/erases) in the backing store.
<b>Total ops (max time/op)</b>	The longest time taken to perform an operation (erase/load/store) in the backing store for any node in the cluster.*
<b>Cluster Trends</b>	Shows the following metrics for the selected cluster. <ul style="list-style-type: none"> <li><b>Events/sec</b> -- Traces the rate of events received in the cluster.</li> <li><b>Rules/ sec</b> -- Traces the rate of rules in the cluster.</li> <li><b>Cache Ops/ sec</b> -- Traces the rate of cache operations in the cluster.</li> <li><b>BS ops/sec</b>-- Traces the rate of backstore operations in the cluster.</li> <li><b>BS max time/op</b>-- Traces the average maximum time per backstore operation.</li> </ul>
<b>Log Scale</b>	This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.

**Base at Zero**

When this option is checked, zero is set as the Y axis minimum for all graph traces.

**Time Range**

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Cluster Nodes Table

Use this display to view configuration and utilization data for nodes in a cluster.

Cluster Name	Node	Alert Severity	Alert Count	Member Count	Auto Startup	Backing Store Enabled	Cache Aside	Serialization Optimized	Storage Enabled	Cache Ty
ckfdcache	new51Cache		0	2	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	DISTRIBUTE
ckfdcache	new51Inf		0	2	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	DISTRIBUTE

**Title Bar** (possible features are):

-   Open the previous and upper display.
-  Open an instance of this display in a new window.
-  Open the online help page for this display.
- open commonly accessed displays.
- The number of items currently in the display.

 **Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

 Open the **Alert Views - RTView Alerts Table** display.

**Filter By:**

The display might include these filtering options:

**Cluster** Choose a cluster for which you want to see metrics.

**Cluster Nodes Table**

Each row in the table is a different node. Data in the row columns describe the node.

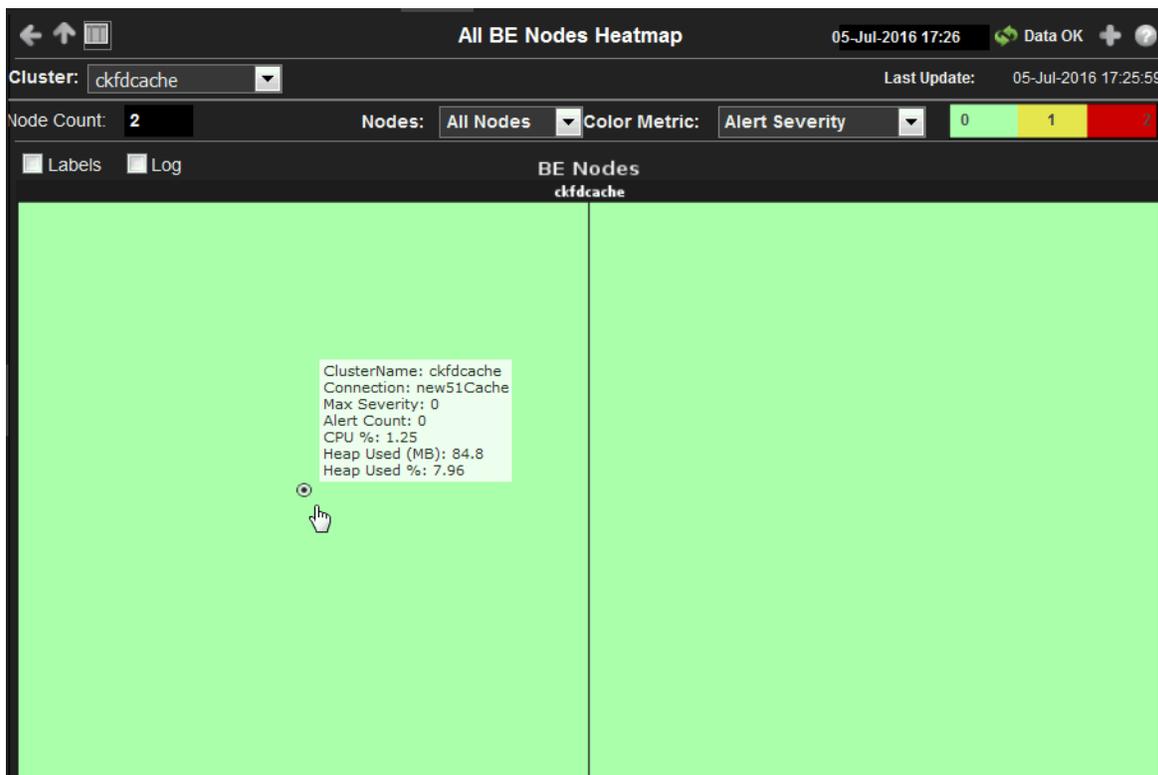
**Note:** Row columns in this table with an asterisk (\*) at the end of the column definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these columns.

<b>Node Count:</b>	The total number of clusters in the table.
<b>Cluster Name</b>	The name of the TIBCO BusinessEvents cluster.
<b>Node</b>	The name of the node.
<b>Alert Severity</b>	The severity level of open alerts. Values range from <b>0</b> to <b>2</b> , where <b>2</b> is the greatest Severity: <ul style="list-style-type: none"> <li> One or more alerts exceeded their ALARM LEVEL threshold.</li> <li> One or more alerts exceeded their WARNING LEVEL threshold.</li> <li> No alert thresholds have been exceeded.</li> </ul>
<b>Alert Count</b>	The total number of critical and warning alerts.
<b>Member Count</b>	The number of neighbors seen by a given node. This value is obtained directly from each node in the cluster. This value should always match the total "Member Count" in the corresponding row of the <b>Clusters</b> table. If they do not match, the node did not join the cluster properly and, hence, the cluster should be restarted.
<b>Auto Startup</b>	When checked ( <b>true</b> ), this feature is enabled.
<b>Backing Store Enabled</b>	When checked ( <b>true</b> ), this feature is enabled.*
<b>Cache Aside</b>	When checked ( <b>true</b> ), this feature is enabled.*
<b>Serialization Optimized</b>	When checked ( <b>true</b> ), this feature is enabled.*
<b>Storage Enabled</b>	When checked ( <b>true</b> ), this feature is enabled.*
<b>Cache Type</b>	The type of TIBCO BusinessEvents cache.*
<b>BE Version</b>	The approximate TIBCO BusinessEvents version, as configured by the <b>connection</b> property. The exact version information is not available via JMX.

<b>Cache Node?</b>	When checked ( <b>true</b> ), the node is a storage node. Otherwise, it is an inference node. This column is added by the Monitor rather than read from the JMX interface.
<b>Node ID</b>	A unique string that identifies the node.
<b>Host</b>	The IP address of the host to which the node is connected.
<b>Port</b>	The port number of the host to which the node is connected.
<b>URL</b>	Uniform Resource Locator, used as an alternative way to specify a JMX connection. When set, the <b>Host</b> and <b>Port</b> columns are blank (and vice versa).
<b>% CPU Used</b>	The amount of CPU, in percent, used by the node. This value is derived from the java.lang.OperatingSystem MBean.
<b>Heap-Max</b>	The maximum amount of memory, in megabytes, that can be used by the JVM for heap space. This value is provided by standard Java MBeans.
<b>Heap-Used</b>	The current amount of memory, in megabytes, in use by the JVM for heap space. This value is provided by standard Java MBeans.
<b>NonHeap Max</b>	The maximum amount of memory, in megabytes, that can be used by the process (not counting heap usage). This value is provided by standard Java MBeans.
<b>NonHeap Used</b>	The current amount of memory, in megabytes, in use by the process (not counting heap usage). This value is provided by standard Java MBeans.
<b>Host OS</b>	The operating system on the host where the node is running.
<b>Connection String</b>	The connection string for the node, which can be the IP address and port of the host that the node is connected to, or the Uniform Resource Locator (which is used as an alternative way to specify a JMX connection).
<b>Connected</b>	When checked ( <b>true</b> ), the node is currently connected to the Data Server via JMX.
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > ( <b>Project Name</b> ) > <b>Solution Package Configuration</b> > <b>TIBCO BusinessEvents</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Timestamp</b>	The date and time, relative to the Data Server, that data was last collected for the node.

## Cluster Nodes Heatmap

This display allows you to view utilization data for all nodes in a cluster in a heatmap format. You can view heatmap data for **All Nodes**, **Inference** nodes, or **Cache** nodes by selecting the desired option from the **Nodes** drop down list. When you click on the heatmap for one of the nodes, the detailed data for that particular node displays in the ["Inference Node Summary"](#) display if you selected an inference node, or in the ["Storage Node Summary"](#) display if you selected a cache node.



### Title Bar (possible features are):

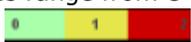
- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- The number of items currently in the display.

- Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

### Filter By:

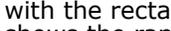
The display might include these filtering options:

- Cluster** Choose a cluster for which you want to see metrics.
- Last Update** The date and time that the display was last updated.
- Node Count** The total number of nodes in the display.

<b>Nodes</b>	Select the type of nodes for which you want to view metrics. You can select from <b>All Nodes</b> , <b>Inference</b> , and <b>Cache</b> . Your selection in this drop down determines the available options in the <b>Color Metric</b> drop down.
<b>Labels</b>	Select this option to display labels in the heatmap for each of the nodes.
<b>Log</b>	Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.
<b>Color Metric</b>	Select the metric driving the heatmap display. The default is Alert Severity. Each <b>Metric</b> has a color gradient bar that maps values to colors. The heatmap organizes the nodes by cluster, where each rectangle represents a node. Mouse-over any rectangle to display the current values of the metrics for the cluster. Click on a rectangle to drill-down to the associated <a href="#">"Storage Node Summary"</a> display for a detailed view of metrics for that particular server. The available options in this drop down change depending on your selection in the <b>Nodes</b> drop down.
<b>Nodes:</b> <b>All Nodes</b>	The following options are available when <b>All Nodes</b> is selected from the <b>Nodes</b> drop down.
<b>Alert Severity</b>	<p>The maximum alert level in the item (index) associated with the rectangle. Values range from <b>0</b> to <b>2</b>, as indicated in the color gradient bar , where <b>2</b> is the greatest <b>Alert Severity</b>.</p> <p><b>2</b> -- Metrics that have exceeded their specified <b>ALARMLEVEL</b> threshold and have an Alert Severity value of <b>2</b> are shown in red. For a given rectangle, this indicates that one or more metrics have exceeded their alarm threshold.</p> <p><b>1</b> -- Metrics that have exceeded their specified <b>WARNINGLEVEL</b> threshold and have an Alert Severity value of <b>1</b> are shown in yellow. For a given rectangle, this indicates that one or more metrics have exceeded their warning threshold.</p> <p><b>0</b> -- Metrics that have not exceeded either specified threshold have an Alert Severity value of <b>0</b> and are shown in green. For a given rectangle, this indicates that no metrics have exceeded a specified alert threshold.</p>
<b>Alert Count</b>	<p>The total number of alarm and warning alerts in a given item (index) associated with the rectangle.</p> <p>The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>JVM % CPU Used</b>	<p>The total percentage of JVM CPU used in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from <b>0</b> to the alert threshold of <b>JvmCpuPercentHigh</b>, which is <b>75</b>. The middle value in the gradient bar indicates the middle value of the range (the default is <b>38</b>).</p>
<b>JVM % Memory Used</b>	<p>The total percentage of JVM Memory Used in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from <b>0</b> to the alert threshold of <b>JvmMemoryUsedHigh</b>, which is <b>75</b>. The middle value in the gradient bar indicates the middle value of the range (the default is <b>38</b>).</p>

**Nodes: Inference** In addition to **Alert Severity, Alert Count, JVM % CPU Used, and JVM % Memory Used**, the following options are also available when **Inference** is selected from the **Nodes** drop down.

**Received Events Rate** The rate of events received in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of **TbeNodeChanRecvdRateHigh**, which is **95**. The middle value in the gradient bar indicates the middle value of the range (the default is **48**).

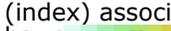
**Rules Fired Rate** The rate of rules fired in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of **TbeNodeRuleFiringRateHigh**, which is **95**. The middle value in the gradient bar indicates the middle value of the range (the default is **48**).

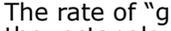
**Total Rules Fired** The total number of rules fired in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the maximum count of rules fired in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

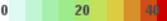
**Nodes: Cache** In addition to **Alert Severity, Alert Count, JVM % CPU Used, and JVM % Memory Used**, the following options are also available when **Cache** is selected from the **Nodes** drop down.

**Backing Store Reads/sec** The rate of reads from the backing store in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of **TbeNodeBackingStoreLoadRateHi**, which is **95**. The middle value in the gradient bar indicates the middle value of the range (the default is **48**).

**Backing Store Writes/sec** The rate of writes to the backing store in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of **TbeNodeBackingStoreStoreRateHi**, which is **95**. The middle value in the gradient bar indicates the middle value of the range (the default is **48**).

**Backing Store Deletes/sec** The rate of deletes from the backing store in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of **TbeNodeBackingStoreEraseRateHigh**, which is **95**. The middle value in the gradient bar indicates the middle value of the range (the default is **48**).

**Concept Gets/sec** The rate of "gets" in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of **TbeNodeConceptsGetRateHigh**, which is **95**. The middle value in the gradient bar indicates the middle value of the range (the default is **48**).

<b>Concept Puts/sec</b>	The rate of “puts” in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from <b>0</b> to the alert threshold of <b>TbeNodeConceptsPutRateHigh</b> , which is <b>95</b> . The middle value in the gradient bar indicates the middle value of the range (the default is <b>48</b> ).
<b>Concept Removes/sec</b>	The rater of “removes” in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from <b>0</b> to the alert threshold of <b>TbeNodeConceptsRemoveRateHigh</b> , which is <b>95</b> . The middle value in the gradient bar indicates the middle value of the range (the default is <b>48</b> ).
<b>Object Table Size</b>	The number of objects maintained in the cache in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from <b>0</b> to the alert threshold of <b>TbeNodeObjectTableSize</b> , which is <b>10,000</b> . The middle value in the gradient bar indicates the middle value of the range (the default is <b>5,000</b> ).

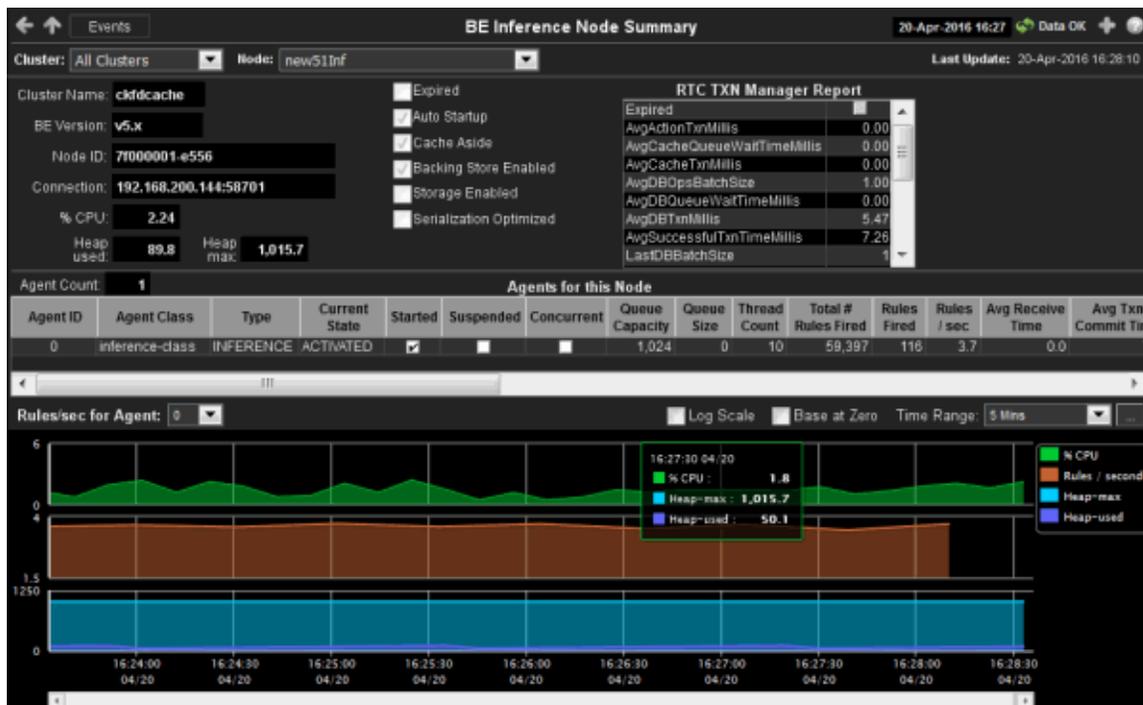
## Inference Node Summary

Use this display to view configuration and utilization data for a single inference node. View a list of all agents on the node, a Run-To-Completion Transaction Manager Report, and trend graphs tracing the rule execution rate for agents on the node. The rule execution rate is relative to the overall CPU and heap utilization for the engine's JVM.

**NOTE:** An inference node (also known as an engine or processing unit) is the container where one or more inference agents run. Generally, the agents in a given node implement different rule sets, and distributing nodes on different hosts provides fault tolerance and load balancing for the cluster. For details, refer to TIBCO documentation.

Choose a single cluster or **All Clusters** and a node from the drop-down menus.

Change the trend graph **Time Range** to “zoom in” on the graph and see more detail or “zoom out” from the graph to see larger trends over time. To change the time range, click Open Time Range , choose the date and time, and then click **OK**.



#### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** **Table** open commonly accessed displays.
- 6,047** The number of items currently in the display.

- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

#### Filter By:

The display might include these filtering options:

- Cluster** Choose a cluster for which you want to view metrics.
- Node** Choose a node for which you want to view metrics.

#### Fields and Data:

- Last Update** The date and time the data in the display was last updated.
- Cluster Name:** The name of the TIBCO BusinessEvents cluster with which the node is a member.
- BE Version:** The approximate TIBCO BusinessEvents version, as configured by the **connection** property. The exact version information is not available via JMX.
- Node ID:** A unique string that identifies the node.

<b>Connection:</b>	The JMX connection method specified in the <b>connection</b> property for a given engine. It is displayed as either a combination of the host and port fields ( <b>&lt;host&gt;:&lt;port&gt;</b> ), or the URL. This convention saves space on the display by avoiding empty fields. This information is provided as a convenience for those rare occasions where a user might wish to view the data directly in jconsole.
<b>% CPU:</b>	The percent of CPU used by the engine process. This value is provided by standard Java MBeans.
<b>Heap used:</b>	The current amount of memory, in megabytes, in use by the JVM for heap space. This value is provided by standard Java MBeans.
<b>Heap max</b>	The maximum amount of memory, in megabytes, that can be used by the JVM for heap space. This value is provided by standard Java MBeans.
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > ( <b>Project Name</b> ) > <b>Solution Package Configuration</b> > <b>TIBCO BusinessEvents</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Auto Startup</b>	When checked ( <b>true</b> ), this feature is enabled.
<b>Cache Aside</b>	When checked ( <b>true</b> ), this feature is enabled.
<b>Backing Store Enabled</b>	When checked ( <b>true</b> ), this feature is enabled.
<b>Storage Enabled</b>	When checked ( <b>true</b> ), this feature is enabled.
<b>Serialization Optimized</b>	When checked ( <b>true</b> ), this feature is enabled.

### RTC TXN Manager Report

**Note:** Fields in this display with an asterisk (\*) at the end of the field definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these fields.

<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > ( <b>Project Name</b> ) > <b>Solution Package Configuration</b> > <b>TIBCO BusinessEvents</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Avg Action Txn Millisec</b>	The average amount of time taken for an action transaction, in milliseconds.*
<b>Avg Cache Queue Wait Time Millisec</b>	The average cache queue wait time, in milliseconds.*
<b>Avg Cache Txn Millisec</b>	The average amount of time taken for a cache transaction, in milliseconds.*
<b>Avg DB Ops Batch Size</b>	The average database operation batch size.*
<b>Avg DB Queue Wait Time Millisec</b>	The average database queue wait time, in milliseconds.*
<b>Avg DB Txn Millisec</b>	The average amount of time taken for a database transaction, in milliseconds.*

<b>Avg Successful Txn Time Millisec</b>	The average amount of time taken for a successful transaction, in milliseconds.*
<b>Last DB Batch Size</b>	The size of the last database batch.*
<b>Pending Actions</b>	The total number of pending actions.*
<b>Pending Cache Writes</b>	The total number of pending cache writes.*
<b>Pending DB Writes</b>	The total number of pending database writes.*
<b>Pending Events to Ack</b>	The total number of pending events that need to be acknowledged.*
<b>Pending Locks to Release</b>	The total number of pending locks that need to be released.*
<b>Total DB Txns Completed</b>	The total number of database transactions that have been completed.*
<b>Total Errors</b>	The total number of errors.*
<b>Total Successful Txns</b>	The total number of successful transactions.*

### Agents for this Node Table

Each row in the table is an agent associated with the node, with data in the row columns describing the agent.

**Note:** Row columns in this table with an asterisk (\*) at the end of the column definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these columns.

<b>Agent Count:</b>	The number of agents currently in the table.
<b>Agent ID</b>	The agent's ID.
<b>Agent Class</b>	The agent's class. See TIBCO documentation for more information.
<b>Type</b>	The type of agent (Inference, Cache, Query, or Dashboard).*
<b>Current State</b>	The current state of the agent.*
<b>Started</b>	When checked, denotes that the agent is started.*
<b>Suspended</b>	When checked, denotes that the agent is suspended.*
<b>Concurrent</b>	When checked, denotes that it is a concurrent agent.*
<b>Queue Capacity</b>	The queue capacity for the agent.*
<b>Queue Size</b>	The queue size for the agent.*
<b>Thread Count</b>	The total number of threads for the agent.*

<b>Total # Rules Fired</b>	The total number of rules fired for the agent.*
<b>Rules Fired</b>	The number of rules fired.*
<b>Rules/sec</b>	The rate of rules fired for the agent.
<b>Avg Receive Time</b>	See TIBCO documentation for more information.*
<b>Avg Txn Commit Time</b>	The average amount of time taken to commit a transaction.*
<b>Cache Queue Remaining</b>	The total amount of remaining space on the cache queue.*
<b>DB Ops Queue Remaining</b>	The total amount of remaining space on the DB Operations queue.*
<b>Hit Ratio</b>	See TIBCO documentation for more information.*
<b>Job Rate</b>	See TIBCO documentation for more information.*
<b>L1 Cache Max Size</b>	The maximum size of the L1 cache.*
<b>L1 Cache Size</b>	The current size of the L1 cache.*
<b>Max Active</b>	See TIBCO documentation for more information.*
<b># Event Threads</b>	The total number of currently active event threads.*
<b># Jobs</b>	The total number of currently active jobs.*
<b>Priority</b>	See TIBCO documentation for more information.*
<b>Read Only</b>	See TIBCO documentation for more information.*
<b>Txn Commit Count</b>	The number of transactions committed by the agent.*
<b>Txn Receive Count</b>	The number of transactions received by the agent.*
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > ( <b>Project Name</b> ) > <b>Solution Package Configuration</b> > <b>TIBCO BusinessEvents</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Timestamp</b>	The date and time, relative to the Data Server, that data was last collected for the agent.

**Trend Graph**

Shows metrics for the selected node.

**% CPU** -- Traces the amount of CPU used, in percent, by the node.

**Rules/sec** -- Traces the number of rules processed, per second, by the agent.

**Heap-max** -- Traces the maximum amount of heap space, in bytes, used by the node since the agent was started.

**Heap-used** -- Traces the current amount of heap space, in bytes, used by the agent.

**Rules/sec for Agent** Choose an agent from the drop-down menu.

**Log Scale**

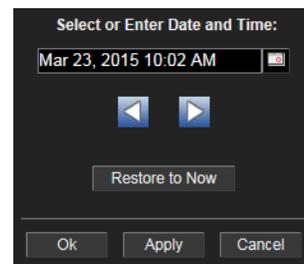
This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.

**Base at Zero**

When this option is checked, zero is set as the Y axis minimum for all graph traces.

**Time Range**

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

**Storage Node Summary**

Use this display to view configuration details for a single cache node, the database connection pool status, as well as a list of all caches that are backed by the backing store (database). Also view trend graphs that trace utilization metrics such as CPU and heap memory usage.

**NOTE:** A storage node (also known as a cache node) provides fast access to events and concepts required during each RTC by the inference engines. Storage nodes also serve as buffers for reads and writes between the cluster and the backing store. For details, refer to TIBCO documentation.

Choose a single cluster or **All Clusters** and a node from the drop-down menus.

Change the trend graph **Time Range** to “zoom in” on the graph and see more detail or “zoom out” from the graph to see larger trends over time. To change the time range click Open Time Range , choose the date and time, then click **OK**.



**Title Bar (possible features are):**

-   Open the previous and upper display.
-  Open an instance of this display in a new window.
-  Open the online help page for this display.
-   open commonly accessed displays.
-  **6,047** The number of items currently in the display.

 **Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

 **23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

 Open the **Alert Views - RTView Alerts Table** display.

**Filter By:**

The display might include these filtering options:

**Cluster:** Choose a cluster to see metrics for.

**Node:** Choose a node to see metrics for.

**Fields and Data**

**Last Update** The date and time the data was last updated in the display.

<b>Cluster Name:</b>	The name of the TIBCO BusinessEvents cluster with which the node is a member.
<b>BE Version:</b>	The approximate TIBCO BusinessEvents version, as configured by the <b>connection</b> property. The exact version information is not available via JMX.
<b>Node ID:</b>	A unique string that identifies the node.
<b>Connection:</b>	The JMX connection method specified in the <b>connection</b> property for a given engine. It is displayed as either a combination of the host and port fields ( <b>&lt;host&gt;:&lt;port&gt;</b> ), or the URL. This convention saves space on the display by avoiding empty fields. This information is provided as a convenience for those rare occasions where a user might wish to view the data directly in jconsole.
<b>% CPU:</b>	The amount of CPU, in percent, used by the node. This value is provided by standard Java MBeans.
<b>Heap used:</b>	The current amount of memory, in megabytes, in use by the JVM for heap space. This value is provided by standard Java MBeans.
<b>Heap Max:</b>	The maximum amount of memory, in megabytes, that can be used by the JVM for heap space. This value is provided by standard Java MBeans.
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > ( <b>Project Name</b> ) > <b>Solution Package Configuration</b> > <b>TIBCO BusinessEvents</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Auto Startup</b>	When checked ( <b>true</b> ), this feature is enabled.
<b>Cache Aside</b>	When checked ( <b>true</b> ), this feature is enabled.
<b>Backing Store Enabled</b>	When checked ( <b>true</b> ), this feature is enabled.
<b>Storage Enabled</b>	When checked ( <b>true</b> ), this feature is enabled.
<b>Serialization Optimized</b>	When checked ( <b>true</b> ), this feature is enabled.

### DB Connection Pool

Values describe status of the pool of database connections used by the cache agent to move data between the local caches and the database.

**Note:** Fields in this region with an asterisk (\*) at the end of the field definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these fields.

<b>Pool State</b>	The state of the database connection pool.*
<b>Auto Failover</b>	The number of times auto failover has occurred.*
<b>Failover Interval</b>	The number of seconds taken for failover to take place.*
<b>Cache Size</b>	The cache size.*
<b># Connections Available</b>	The total number of connections available.*
<b># Connections in Use</b>	The total number of connections currently in use.*

**Backing StoreTable**

A cache node manages access to current events and concepts, buffering as necessary between local memory and a database. The Backing Store table provides a list of caches and the database select/insert/delete statistics for each cache.

**Note:** Row columns in this table with an asterisk (\*) at the end of the column definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these columns.

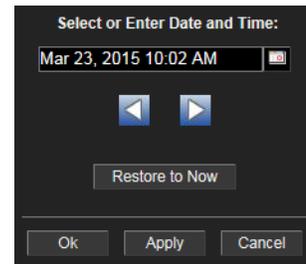
<b>Cache Name</b>	The name of the cache.*
<b>Active</b>	When checked, denotes that the cache is active.*
<b>Delete Avg Time</b>	The average amount of time taken for a "delete" ("erase") operation in the Backing Store for the cache.*
<b>Load Avg Time</b>	The average amount of time taken for a "load" operation in the Backing Store for the cache.*
<b>Store Avg Time</b>	The average amount of time taken for a "store" operation in the Backing Store for the cache.*
<b>Delete Total</b>	The total number of "delete" operations in the Backing Store for the cache.*
<b>Load Total</b>	The total number of "load" operations in the Backing Store for the cache.*
<b>Store Total</b>	The total number of "store" operations in the Backing Store for the cache.*
<b>Deletes</b>	The number of "delete" operations during the last polling interval.*
<b>Loads</b>	The number of "load" operations during the last polling interval.*
<b>Stores</b>	The number of "store" operations during the last polling interval.*
<b>Deletes/sec</b>	The rate of "delete" operations in the node.
<b>Loads/sec</b>	The rate of "load" operations in the node.
<b>Stores/sec</b>	The rate of "store" operations in the node.
<b>Object Table Trends</b>	Shows metrics for the selected cluster/node combination: <ul style="list-style-type: none"> <li><b>% CPU</b> -- Traces the amount of CPU used, in percent, by the engine.</li> <li><b>Table Size</b> -- Traces the number of unique objects cached in the local index table.</li> <li><b>Ext ID Tbl Size</b> -- Traces the number of entries in the table of external IDs used as indexes by the backing store.</li> <li><b>Max Heap (MB)</b> -- Traces the maximum amount of memory, in megabytes, that can be used by the JVM for heap space.</li> <li><b>Heap (MB)</b> -- Traces the current heap space, in megabytes, in use by the JVM.</li> </ul>
<b>Log Scale</b>	This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.

**Base at Zero**

When this option is checked, zero is set as the Y axis minimum for all graph traces.

**Time Range**

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows  to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Events / Concepts View

These displays present performance data for your BusinessEvents system. Displays in this View are:

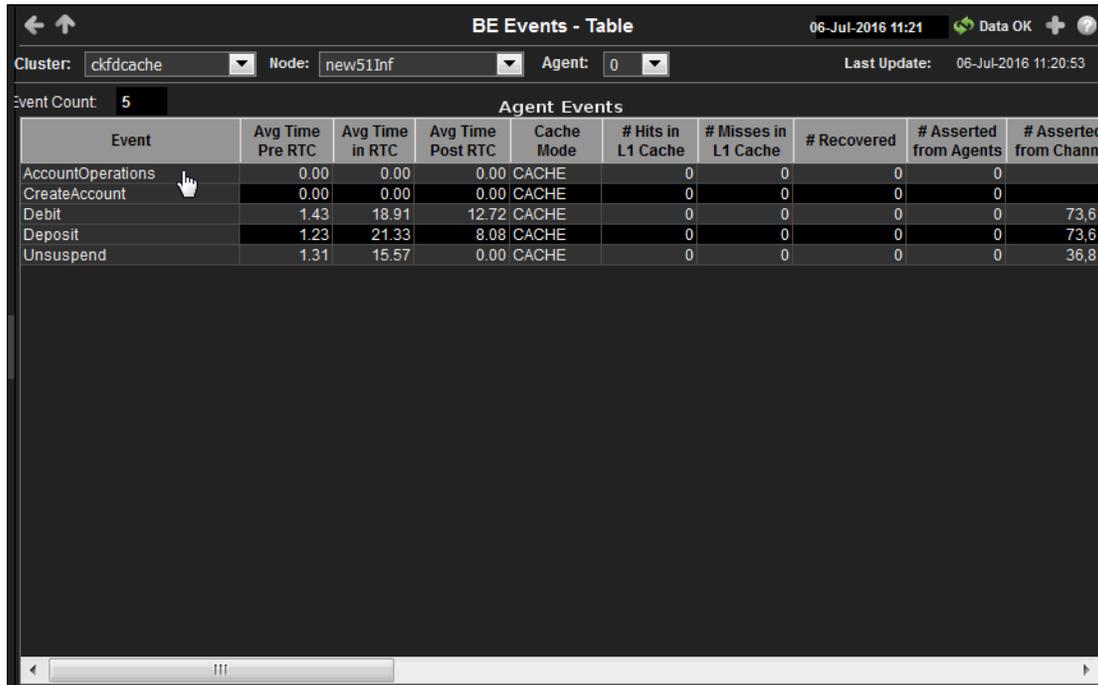
- "Agent Events"
- "Agent Event Summary"
- "Event Cache Hits"
- "Event Hit Summary"
- "Concept Cache Hits"
- "Concept Hit Summary"
- "Channels"
- "All Inference Agents"
- "All RTC Reports"

### Agent Events

View run-time statistics for a selected group of agents. With TIBCO BusinessEvents, events are cached when they are out of scope, and deleted or persisted to the backing store when they are no longer useful. Clicking on a row in the table displays access patterns over time for the event in the "Agent Event Summary" display.

**NOTE:** Events cause rules to execute in the BusinessEvents Rete network. Events can be created by external phenomena, such as the arrival of a JMS message, or internally when rules are processed. When an event enters the Rete network, it causes a run-to-completion cycle which continues until no further rules can be processed. Each named event that can be handled by a BusinessEvents application is specified at build time in BusinessEvents studio. For details, refer to TIBCO documentation.

Sort  the table columns when all the rows cannot fit on the screen. For example, sort  the **Expired** column so that all expired nodes are listed at the top.



Event	Avg Time Pre RTC	Avg Time in RTC	Avg Time Post RTC	Cache Mode	# Hits in L1 Cache	# Misses in L1 Cache	# Recovered	# Asserted from Agents	# Asserted from Chann
AccountOperations	0.00	0.00	0.00	CACHE	0	0	0	0	
CreateAccount	0.00	0.00	0.00	CACHE	0	0	0	0	
Debit	1.43	18.91	12.72	CACHE	0	0	0	0	73.6
Deposit	1.23	21.33	8.08	CACHE	0	0	0	0	73.6
Unsuspend	1.31	15.57	0.00	CACHE	0	0	0	0	36.8

**Title Bar (possible features are):**

-   Open the previous and upper display.
-  Open an instance of this display in a new window.
-  Open the online help page for this display.
- open commonly accessed displays.
- The number of items currently in the display.
-  **Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
-  Open the **Alert Views - RTView Alerts Table** display.

**Filter By:**

The display might include these filtering options:

- Cluster:** Select the cluster containing the node and agent for which you want to view metrics.
- Node:** Select a node containing the agent for which you want to view metrics.
- Agent** Select the agent for which you want to view metrics.

**Fields and Data:**

**Last Update:** The date and time the data on the display was last updated.

### Agent Events Table:

Each row in the table is a different event. Data in the row columns describe the event. The following fields are added by Monitor collection. The assertions/sec, modified/sec, and retracted/sec metrics are calculated from the corresponding counters as the delta between two successive samples divided by the polling interval.

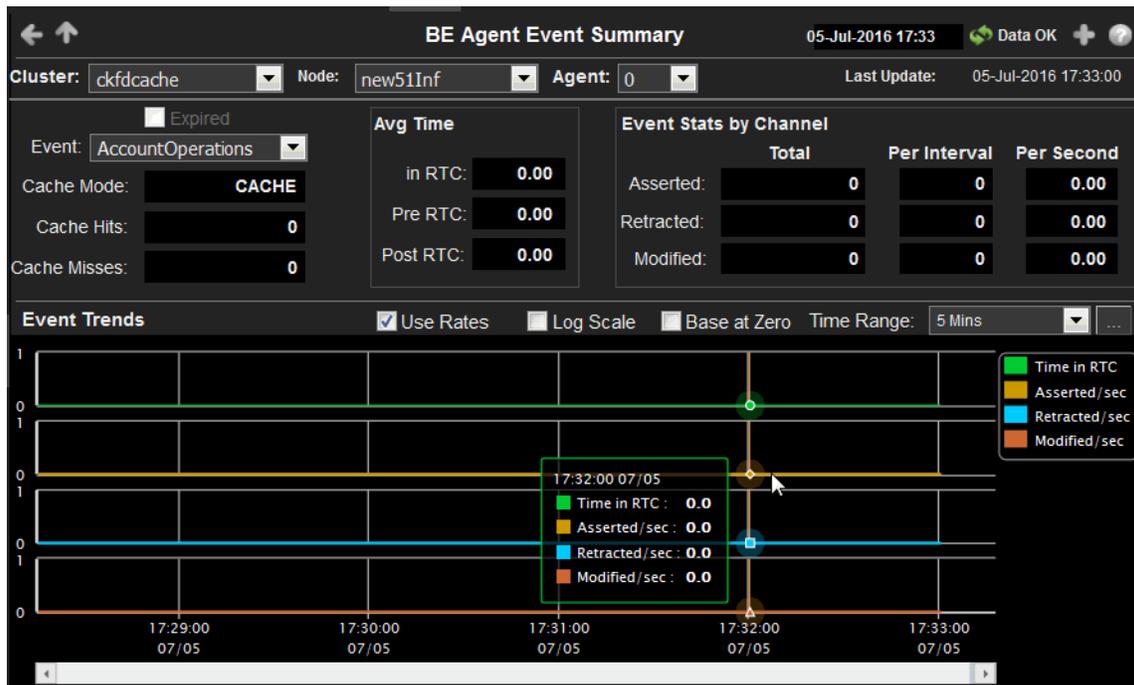
**Note:** Row columns in this table with an asterisk (\*) at the end of the column definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these columns.

<b>Event Count:</b>	The total number of events in the table.
<b>Event</b>	The name of the event.
<b>Avg Time Pre RTC</b>	The average amount of time taken for the event to begin its run to completion cycle.*
<b>Avg Time in RTC</b>	The average amount of time taken for the event to complete (once it has started) its run to completion cycle.*
<b>Avg Time Post RTC</b>	The average amount of time taken by the event after its run to completion cycle has ended.*
<b>Cache Mode</b>	Lists the mode used by the event, which can be either <b>CACHE</b> (only) or <b>MEMORY</b> (only).*
<b># Hits in L1 Cache</b>	The number of times data has been searched for in the L1 cache since the last data update.*
<b># Misses n L1 Cache</b>	The number of times data has been searched for in the L1 cache, but was not found, since the last data update.*
<b># Recovered</b>	The number of times data is not found in the L1 cache, but is found in a different cache, since the last data update.*
<b># Asserted from Agents</b>	The number of times the event was asserted by an agent into the Rete network.*
<b># Asserted from Channel</b>	The number of times the event was asserted into the Rete network via the channel.*
<b># Modified from Agents</b>	The number of times the event was modified by an agent in the Rete network.*
<b># Modified from Channel</b>	The number of times the event was modified in the Rete network via the channel.*
<b># Retracted from Agents</b>	The number of times the event was retracted/deleted by an agent from the Rete network.*
<b># Retracted from Channel</b>	The number of times the event was retracted/deleted from the Rete network via the channel.*
<b>L1 Cache Hits/sec</b>	The rate of L1 cache hits.
<b>L1 Cache Misses/sec</b>	The rate of L1 cache misses.
<b># Recovered / sec</b>	The rate of recovered data.

- Assertions/ sec (Agent)**     The rate of event assertions into the Rete network by the agent.
- Assertions/ sec (Channel)**     The rate of event assertions into the Rete network via the channel.
- Modifies/ sec (Agent)**     The rate of event modifications in the Rete network by the agent.
- Modifies/ sec (Channel)**     The rate of event modifications in the Rete network via the channel.
- Retractions /sec (Agent)**     The rate of event retractions/deletions from the Rete network by the agent.
- Retractions /sec (Channel)**     The rate of event retractions/deletions from the Rete network via the channel.
- Expired**     When checked, performance data has not been received within the time specified (in seconds) in the **Expire Time** field in the **Duration** region in the RTView Configuration Application > **(Project Name)** > **Solution Package Configuration** > **TIBCO BusinessEvents** > **DATA STORAGE** tab. The **Delete Time** field (also in the **Duration** region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
- Timestamp**     The date and time, relative to the Data Server, that data was last collected for the engine.

### Agent Event Summary

View detailed performance metrics for an agent’s event. You can view cache, RTC, event statistics by channel, and event trend data over a specified period of time.



**Title Bar** (possible features are):

-   Open the previous and upper display.
-  Open an instance of this display in a new window.
-  Open the online help page for this display.
- Menu**  **Table** open commonly accessed displays.
- 6,047** The number of items currently in the display.

 **Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

 Open the **Alert Views - RTView Alerts Table** display.

**Filter By:**

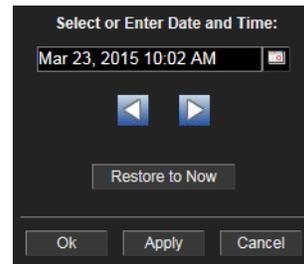
The display might include these filtering options:

**Note:** Fields in this display with an asterisk (\*) at the end of the column definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these fields.

<b>Cluster:</b>	Select the cluster for which you want to see metrics.						
<b>Node:</b>	Select the node for which you want to see metrics.						
<b>Agent</b>	Select the agent for which you want to see metrics.						
<b>Last Update</b>	The date and time in which the data was last updated.						
<b>Expired</b>	When checked ( <b>true</b> ), the Monitor has not received a response from the event for the amount of time specified by the <b>\$tbeRowExpirationTime</b> property (the default is <b>120</b> seconds). When the amount of time specified by the <b>\$tbeRowExpirationTimeForDelete</b> property elapses (the default is one day), the event data is deleted from the cache and display.						
<b>Event</b>	The name of the event.						
<b>Cache Mode</b>	Lists the mode used by the event, which can be either <b>CACHE</b> (only) or <b>MEMORY</b> (only).*						
<b>Cache Hits</b>	The number of times data has been searched for in the L1 cache since the last data update.*						
<b>Cache Misses</b>	The number of times data has been searched for in the L1 cache, but was not found, since the last data update.*						
<b>Avg Time</b>	<table border="0"> <tr> <td><b>in RTC</b></td> <td>The average amount of time taken for the event to complete (once it has started) its run to completion cycle.*</td> </tr> <tr> <td><b>Pre RTC</b></td> <td>The average amount of time taken for the event to begin its run to completion cycle.*</td> </tr> <tr> <td><b>Post RTC</b></td> <td>The average amount of time taken by the event after its run to completion cycle has ended.*</td> </tr> </table>	<b>in RTC</b>	The average amount of time taken for the event to complete (once it has started) its run to completion cycle.*	<b>Pre RTC</b>	The average amount of time taken for the event to begin its run to completion cycle.*	<b>Post RTC</b>	The average amount of time taken by the event after its run to completion cycle has ended.*
<b>in RTC</b>	The average amount of time taken for the event to complete (once it has started) its run to completion cycle.*						
<b>Pre RTC</b>	The average amount of time taken for the event to begin its run to completion cycle.*						
<b>Post RTC</b>	The average amount of time taken by the event after its run to completion cycle has ended.*						
<b>Event Stats by Channel</b>	<table border="0"> <tr> <td><b>Asserted Total</b></td> <td>The total number of times the event was asserted into the Rete network via the channel.*</td> </tr> <tr> <td><b>Asserted Per Interval</b></td> <td>The number of times the event was asserted into the Rete network via the channel since the last data update.*</td> </tr> <tr> <td><b>Asserted Per Second</b></td> <td>The rate of event assertions into the Rete network via the channel.</td> </tr> </table>	<b>Asserted Total</b>	The total number of times the event was asserted into the Rete network via the channel.*	<b>Asserted Per Interval</b>	The number of times the event was asserted into the Rete network via the channel since the last data update.*	<b>Asserted Per Second</b>	The rate of event assertions into the Rete network via the channel.
<b>Asserted Total</b>	The total number of times the event was asserted into the Rete network via the channel.*						
<b>Asserted Per Interval</b>	The number of times the event was asserted into the Rete network via the channel since the last data update.*						
<b>Asserted Per Second</b>	The rate of event assertions into the Rete network via the channel.						

<b>Retracted Total</b>	The total number of times the event was retracted/deleted from the Rete network via the channel.*
<b>Retracted Per Interval</b>	The number of event retractions/deletions from the Rete network.
<b>Retracted Per Second</b>	The rate of event retractions/deletions from the Rete network via the channel.
<b>Modified Total</b>	The total number of times the event was modified in the Rete network via the channel.*
<b>Modified Per Interval</b>	The number of event modifications in the Rete network via the channel.
<b>Modified Per Second</b>	The rate of event modifications in the Rete network via the channel.
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > <b>(Project Name)</b> > <b>Solution Package Configuration</b> > <b>TIBCO BusinessEvents</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Event Trends</b>	Shows metrics for the selected event: <ul style="list-style-type: none"> <li><b>Time in RTC</b>-- Traces the event spends in the run to completion cycle.</li> <li><b>Asserted(/sec)</b>-- Traces the number of events asserted into the Rete network (or the rate of event assertions per second depending on <b>Use Rates</b> setting).</li> <li><b>Retracted(/sec)</b>-- Traces the number events retracted from the Rete network (or rate of event retractions per second depending on <b>Use Rates</b> setting).</li> <li><b>Modified(/sec)</b>-- Traces the number of events modified in the Rete network (or rate of events modified per second depending on <b>Use Rates</b> setting).</li> </ul>
<b>Use Rates</b>	When selected, this toggle allows you to view data in the trend graph in counts per second (asserted count per second, retracted count per second, and modified count per second) instead of the default counts per selected interval (asserted count, retracted count, modified count).
<b>Log Scale</b>	This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.

- Base at Zero** When this option is checked, zero is set as the Y axis minimum for all graph traces.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows  to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

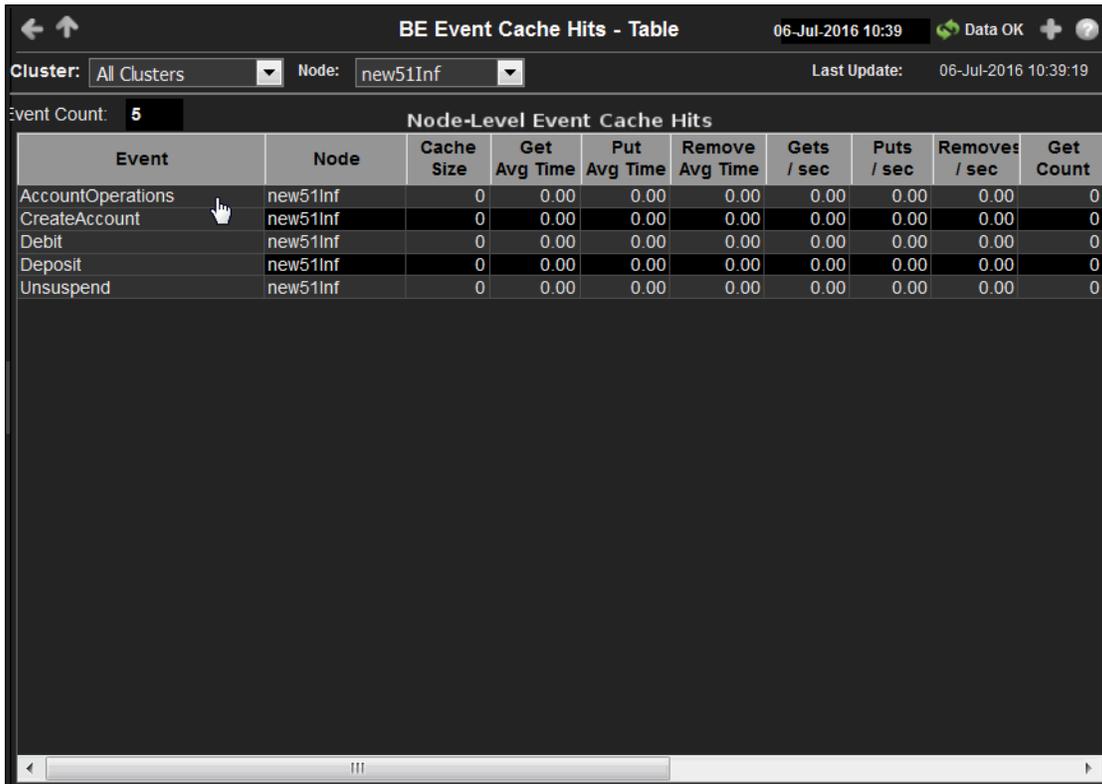
## Event Cache Hits

View cache performance metrics per event for a single cluster or **All Clusters**.

**NOTE:** Events cause rules to execute in the BusinessEvents Rete network. Events can be created by external phenomena, such as the arrival of a JMS message, or internally when rules are processed. When an event enters the Rete network, it causes a run-to-completion cycle which continues until no further rules can be processed. Each named event that can be handled by a BusinessEvents application is specified at build time in BusinessEvents studio. For details, refer to TIBCO documentation.

Choose a single cluster or **All Clusters** and a node from the drop-down menus.

Sort  the table columns when all the rows cannot fit on the screen. For example, sort  the **Expired** column so that all expired nodes are listed at the top.



Event	Node	Cache Size	Get Avg Time	Put Avg Time	Remove Avg Time	Gets / sec	Puts / sec	Removes / sec	Get Count
AccountOperations	new51Inf	0	0.00	0.00	0.00	0.00	0.00	0.00	0
CreateAccount	new51Inf	0	0.00	0.00	0.00	0.00	0.00	0.00	0
Debit	new51Inf	0	0.00	0.00	0.00	0.00	0.00	0.00	0
Deposit	new51Inf	0	0.00	0.00	0.00	0.00	0.00	0.00	0
Unsuspend	new51Inf	0	0.00	0.00	0.00	0.00	0.00	0.00	0

**Title Bar** (possible features are):

-   Open the previous and upper display.
-  Open an instance of this display in a new window.
-  Open the online help page for this display.
- Menu** , **Table** open commonly accessed displays.
- 6,047** The number of items currently in the display.
-  **Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
-  Open the **Alert Views - RTView Alerts Table** display.

**Filter By:**

The display might include these filtering options:

- Cluster:** Select a cluster for which you want to see metrics.
- Node:** Select a node for which you want to see metrics.
- Last Update** The date and time the data was last updated.

**Node-Level Event Statistics Table:**

Each row in the table is a different event, with data in the row columns describing the event.

**Note:** Row columns in this table with an asterisk (\*) at the end of the column definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these columns.

<b>Event Count:</b>	The total number of events in the table.
<b>Event</b>	The name of the event.
<b>Node</b>	The name of the node.
<b>Cache Size</b>	The size of the event's cache.*
<b>Get Avg Time</b>	The average time taken for a "get" event for the node.*
<b>Put Avg Time</b>	The average time taken for a "put" event for the node.*
<b>Remove Avg Time</b>	The average time taken for a "remove" event for the node.*
<b>Gets/sec</b>	The rate of "get" operations for the event.
<b>Puts/sec</b>	The rate of "put" operations for the event.
<b>Removes/sec</b>	The rate of "remove" operations for the event.
<b>Get Count</b>	The total number of "get" operations for the event.*
<b>Put Count</b>	The total number of "put" operations for the event.*
<b>Remove Count</b>	The total number of "remove" operations for the event.*
<b>Num Handles In Store</b>	The number of handles in the Backing Store for the event.*
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > ( <b>Project Name</b> ) > <b>Solution Package Configuration</b> > <b>TIBCO BusinessEvents</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Timestamp</b>	The date and time, relative to the Data Server, that data was last collected for the engine.

## Event Hit Summary

View detailed event performance metrics for a single cluster or **All Clusters**, a node, and an event.

**NOTE:** Events cause rules to execute in the BusinessEvents Rete network. Events can be created by external phenomena, such as the arrival of a JMS message, or internally when rules are processed. When an event enters the Rete network, it causes a run-to-completion cycle which continues until no further rules can be processed. Each named event that can be handled by a BusinessEvents application is specified at build time in BusinessEvents studio. For details, refer to TIBCO documentation.

Choose a single cluster or **All Clusters**, a node and an event from the drop-down menus.



**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** , **Table** open commonly accessed displays.
- 6,047** The number of items currently in the display.
- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

**Filter By:**

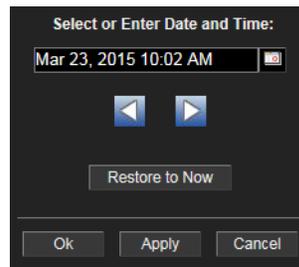
The display might include these filtering options:

**Note:** Fields in this display with an asterisk (\*) at the end of the field definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these fields.

- Cluster:** Select a cluster containing the node and event for which you want to see metrics.
- Node:** Select a node containing the event for which you want to see metrics.
- Event** Select the event for which you want to see metrics.
- Last Update** The date and time in which the data was last updated.

<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > ( <b>Project Name</b> ) > <b>Solution Package Configuration</b> > <b>TIBCO BusinessEvents</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Cache Size</b>	The size of the cache.*
<b>Handles in Store</b>	The number of handles in the Backing Store for the event.*
<b>Avg Cache Access Time</b>	<p><b>Get</b>            The average time taken for a "get" operation.*</p> <p><b>Put</b>             The average time taken for a "put" operation.*</p> <p><b>Remove</b>        The average time taken for a "remove" operation.*</p>
<b>Cache Access Stats</b>	<p><b>Get -- Total Hits</b>     The total number of "get" operations for the event.*</p> <p><b>Get-- Hits</b>         The number of "get" operations for the event since the last data update.*</p> <p><b>Get-- Hits/sec</b>     The rate of "get" operations for the event.</p> <p><b>Put-- Total Hits</b>     The total number of "put" operations for the event.*</p> <p><b>Put--Hits</b>         The number of "put" operations for the event since the last data update.*</p> <p><b>Put-- Hits/sec</b>     The rate of "put" operations for the event.</p> <p><b>Remove --Total Hits</b>     The total number of "remove" operations for the event.*</p> <p><b>Remove--Hits</b>     The number of "remove" operations for the event since the last data update.*</p> <p><b>Remove --Hits/sec</b>     The rate of "remove" operations for the event.</p>
<b>Cache Access Trends</b>	<p>Shows metrics for the selected cluster/node/event combination:</p> <p><b>Gets(/sec)</b> -- Traces the number of "gets" (or rate of "gets" per second depending on <b>Use Rates</b> setting) for the event.</p> <p><b>Puts(/sec)</b>-- Traces the number of "puts" (or rate of "puts" per second depending on <b>Use Rates</b> setting) for the event.</p> <p><b>Removes(/sec)</b>-- Traces the number of "removes" (or rate of "removes" per second depending on <b>Use Rates</b> setting) for the event.</p> <p><b>Use Rates</b>            When selected, this toggle allows you to view data in the trend graph in counts per second ("get" operations count per second, "put" operations count per second, and "remove" operations count per second) instead of the default counts per selected interval ("get" operations count, "put" operations count, "remove" operations count).</p>

- Log Scale** This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.
- Base at Zero** When this option is checked, zero is set as the Y axis minimum for all graph traces.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

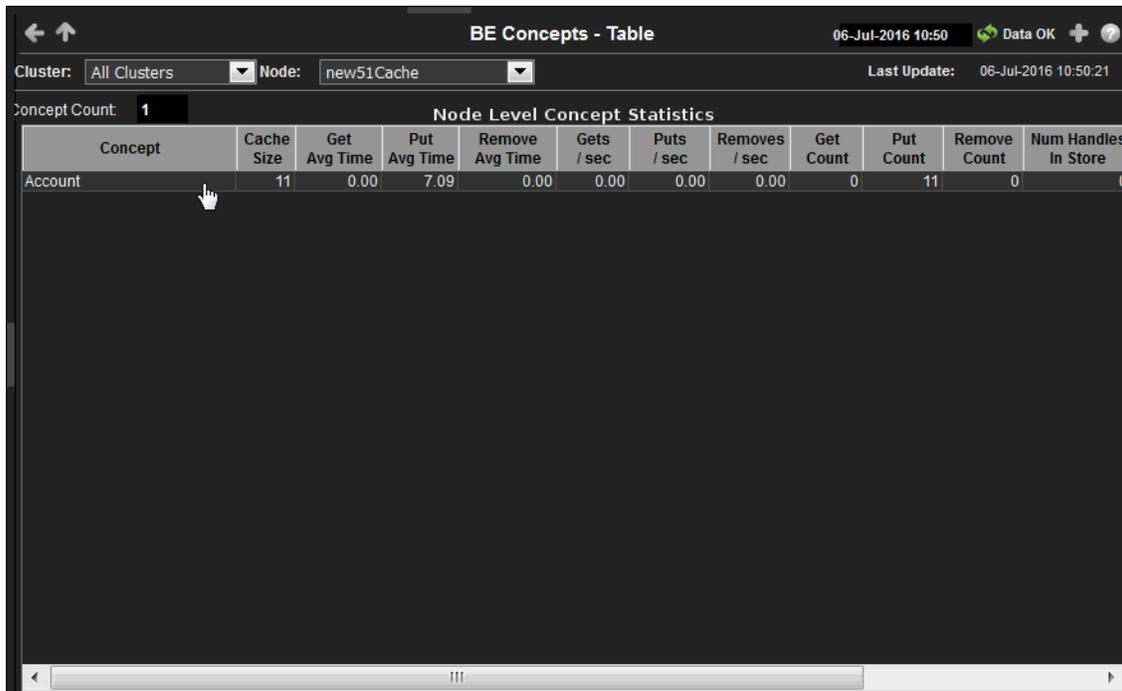
Use the navigation arrows   to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Concept Cache Hits

View a list of concepts and their run-time statistics. Choose a single cluster or **All Clusters** and a node from the drop-down menus.

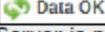
Sort  the table columns when all the rows cannot fit on the screen. For example, sort  the **Expired** column so that all expired nodes are listed at the top.



Concept	Cache Size	Get Avg Time	Put Avg Time	Remove Avg Time	Gets / sec	Puts / sec	Removes / sec	Get Count	Put Count	Remove Count	Num Handles In Store
Account	11	0.00	7.09	0.00	0.00	0.00	0.00	0	11	0	0

### Title Bar (possible features are):

-   Open the previous and upper display.
-  Open an instance of this display in a new window.
-  Open the online help page for this display.
- Menu**  **Table** open commonly accessed displays.
- 6,047** The number of items currently in the display.

-  **Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
-  Open the **Alert Views - RTView Alerts Table** display.

### Filter By:

The display might include these filtering options:

- Cluster:** Choose a cluster to see metrics for.
- Node:** Choose a node to see metrics for.
- Last Update** The date and time the data was last updated.

**Node-Level Concept Statistics Table:**

Each row in the table provides statistics regarding data access for a given BusinessEvents concept.

**Note:** Row columns in this table with an asterisk (\*) at the end of the column definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these columns.

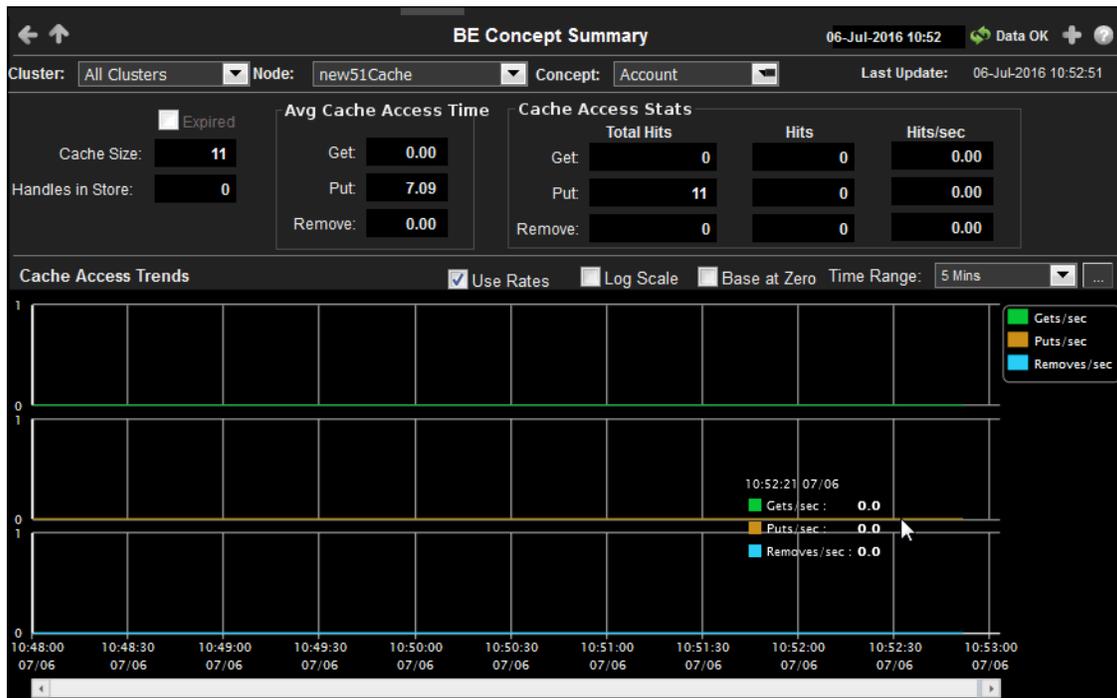
<b>Concept Count:</b>	The total number of concepts in the table.
<b>Concept</b>	The name of the concept.
<b>Cache Size</b>	The size of the concept's cache.*
<b>Get Avg Time</b>	The average time taken for a "get" operation.*
<b>Put Avg Time</b>	The average time taken for a "put" operation.*
<b>Remove Avg Time</b>	The average time taken for a "remove" operation.*
<b>Gets/sec</b>	The rate of "gets" for the concept.
<b>Puts/sec</b>	The rates of "puts" for the concept.
<b>Removes/sec</b>	The rate of "removes" for the concept.
<b>Get Count</b>	The total number of "gets" for the concept.*
<b>Put Count</b>	The total number of "puts" for the concept.*
<b>Remove Count</b>	The total number of "removes" for the concept.*
<b>Num Handles In Store</b>	The number of handles in the Backing Store for the concept.*
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > ( <b>Project Name</b> ) > <b>Solution Package Configuration</b> > <b>TIBCO BusinessEvents</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Timestamp</b>	The date and time, relative to the Data Server, that data was last collected for the concept.

**Concept Hit Summary**

Use this display to view current and historic data for a single concept. Data in this display can be useful if your BusinessEvents system uses Cache object management. When Cache object management is used, concepts with a sufficiently long time to live (TTL) setting are cached.

Cache reference patterns for certain concepts may be related to incoming events (for example, customer purchase orders with associated inventory queries). The trend charts show the cache activity of such concepts, and might be useful in diagnosing the behavior of your application over time.

Choose a single cluster or **All Clusters**, a node and a concept from the drop-down menus. Change the trend graph **Time Range** to "zoom in" on the graph and see more detail or "zoom out" from the graph to see larger trends over time.



#### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- The number of items currently in the display.

- Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

#### Filter By:

Fields in this table with an asterisk (\*) at the end of the field definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these fields. The display might include these filtering options:

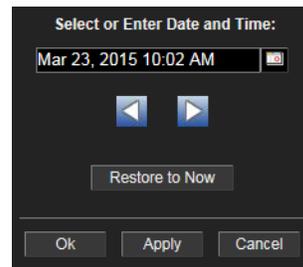
- Cluster:** Select a cluster containing the node and concept for which you want to see metrics.
- Node:** Select a node containing the concept for which you want to see metrics.
- Concept:** Select the concept for which you want to see metrics.

#### Fields and Data:

**Note:** Fields in this table with an asterisk (\*) at the end of the field definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these fields.

<b>Last Update</b>	The date and time in which the data was last updated in the display.
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > ( <b>Project Name</b> ) > <b>Solution Package Configuration</b> > <b>TIBCO BusinessEvents</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Cache Size</b>	The size of the cache.*
<b>Handles in Store</b>	The number of handles in the Backing Store.*
<b>Avg Cache Access Time</b>	<p><b>Get</b>            The average time taken for a "get" operation.*</p> <p><b>Put</b>             The average time taken for a "put" operation.*</p> <p><b>Remove</b>        The average time taken for a "remove" operation.*</p>
<b>Cache Access Stats</b>	<p><b>Get -- Total Hits</b>        The total number of "get" operations for the concept.*</p> <p><b>Get-- Hits</b>            The number of "get" operations for the concept since the last data update.*</p> <p><b>Get-- Hits/sec</b>        The rate of "get" operations for the concept.</p> <p><b>Put-- Total Hits</b>        The total number of "put" operations for the concept.*</p> <p><b>Put--Hits</b>            The number of "put" operations for the concept since the last data update.*</p> <p><b>Put-- Hits/sec</b>        The rate of "put" operations for the concept.</p> <p><b>Remove --Total Hits</b>    The total number of "remove" operations for the concept.*</p> <p><b>Remove- Hits</b>         The number of "remove" operations for the concept since the last data update.*</p> <p><b>Remove --Hits/sec</b>    The rate of "remove" operations for the concept.</p>
<b>Cache Access Trends</b>	Shows metrics for the selected cluster/node/concept combination: <p><b>Gets(/sec)</b> -- Traces the number of "get" operations (or rate of "get" operations depending on <b>Use Rates</b> setting) for the concept.</p> <p><b>Puts(/sec)</b>-- Traces the number of "put" operations (or rate of "put" operations depending on <b>Use Rates</b> setting) for the concept.</p> <p><b>Removes(/sec)</b>-- Traces the number of "remove" operations (or rate of "remove" operations depending on <b>Use Rates</b> setting) for the concept.</p>

- Use Rates** When selected, this toggle allows you to view data in the trend graph in counts per second ("get" operations count per second, "put" operations count per second, and "remove" operations count per second) instead of the default counts per selected interval ("get" operations count, "put" operations count, "remove" operations count).
- Log Scale** This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.
- Base at Zero** When this option is checked, zero is set as the Y axis minimum for all graph traces.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

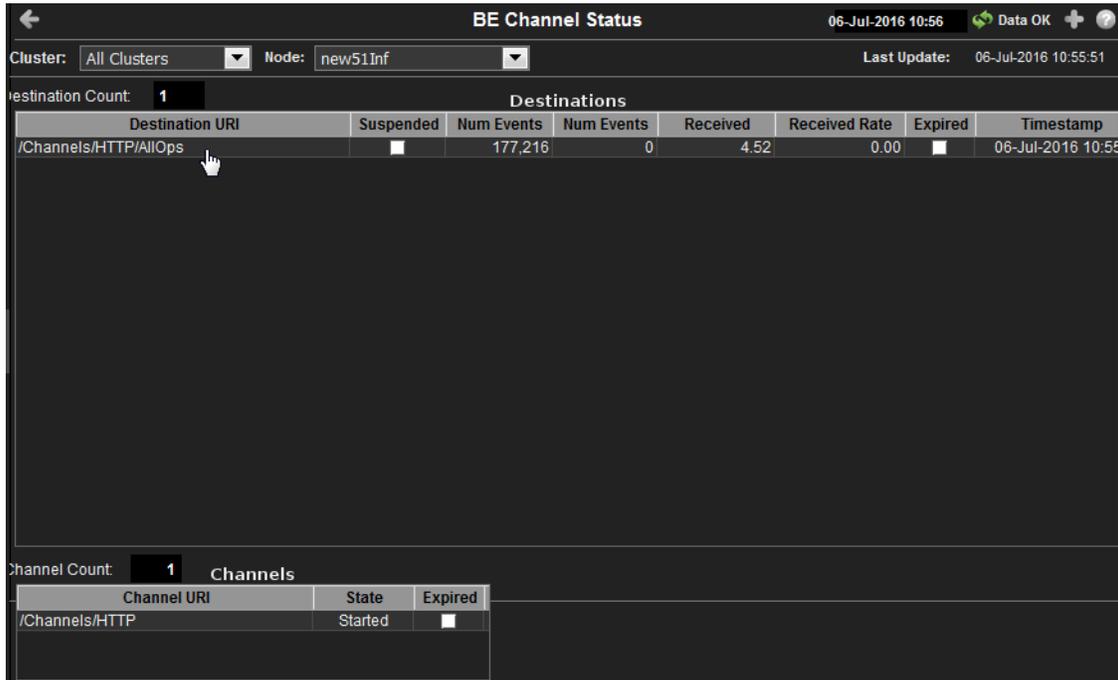
Click **Restore to Now** to reset the time range end point to the current time.

## Channels

Use this display to view a list of destinations, which are sources and sinks of events. Destinations are potentially bi-directional, and the table indicates whether events are sent or received.

**NOTE:** Channels provide a class wrapper for destinations, and make it possible to enable or disable a group of destinations with one operation.

Choose a single cluster or **All Clusters** and a node from the drop-down menus. Each row in the table is a different destination URI. Click a row to view channel details in the **Channels** table.



**Title Bar (possible features are):**

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu, Table open commonly accessed displays.
- 6,047 The number of items currently in the display.
- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

**Filter By:**

The display might include these filtering options:

- Cluster:** Choose a cluster to see metrics for.
- Node:** Choose a node to see metrics for.

**Destinations Table**

Each row in the table provides data for a particular destination.

**Note:** Row columns in this table with an asterisk (\*) at the end of the column definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these columns.

- Destination Count:** The total number of destinations in the table.
- Destination URI:** The Uniform Resource Identifier (URI) for the destination.\*

<b>Suspended</b>	Denotes whether the destination is suspended.*
<b>Num Events Received</b>	The number of events received by the destination.*
<b>Number of Events Sent</b>	The number of events sent by the destination.*
<b>Received Events Rate</b>	The rate of events received by the destination.
<b>Received Rate Last Interval</b>	The rate of events received.
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > ( <b>Project Name</b> ) > <b>Solution Package Configuration</b> > <b>TIBCO BusinessEvents</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Timestamp</b>	The date and time, relative to the Data Server, that data was last collected for the destination.

### Channels Table

Each row in the table provides data for a particular channel.

**Note:** Row columns in this table with an asterisk (\*) at the end of the column definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these columns.

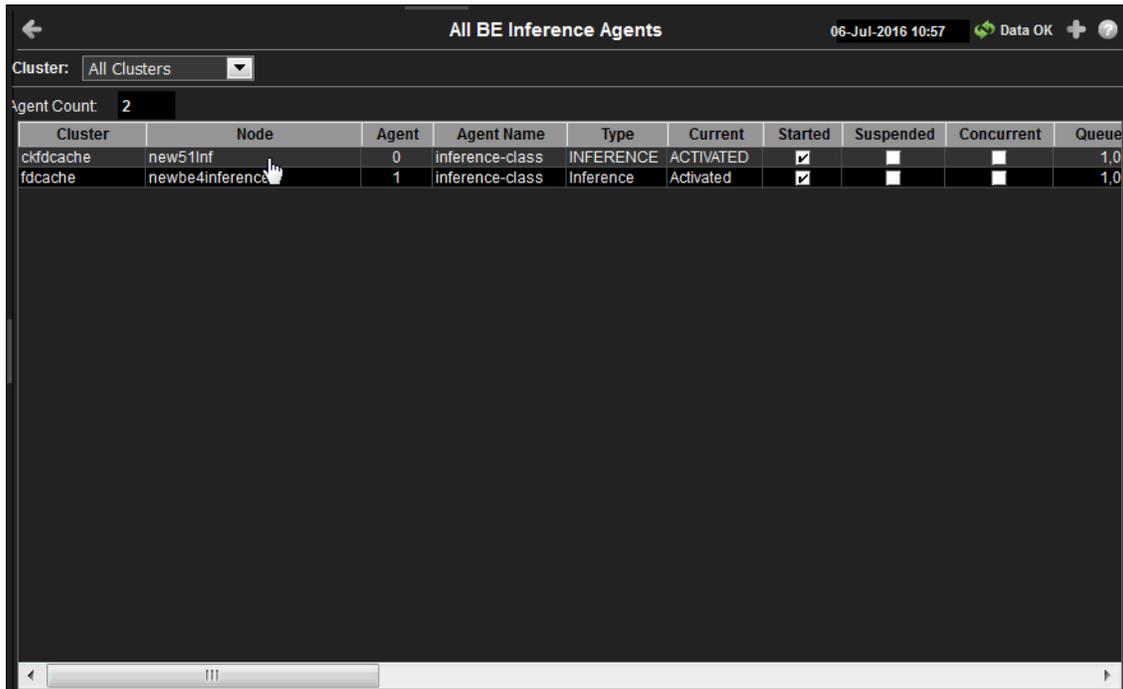
<b>Channel Count:</b>	The total number of channels in the table.
<b>Channel URI</b>	The Uniform Resource Identifier (URI) for the channel.*
<b>State</b>	The current state of the channel.*
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > ( <b>Project Name</b> ) > <b>Solution Package Configuration</b> > <b>TIBCO BusinessEvents</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.

## All Inference Agents

Use this display to compare agent metrics across deployed engines and verify that the cluster is properly load-balanced. View a list of all the inference agents deployed in each cluster. You can view agent data for a single cluster or all clusters.

The data in this display is identical to the data provided for a single engine in the "[Cluster Summary](#)" display, except that it is aggregated across all inference nodes.

Choose a single cluster or **All Clusters** from the drop-down menus. Each row in the table is a different agent.



**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- 6,047** The number of items currently in the display.
- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

**Filter By:**

The display might include these filtering options:

**Cluster:** Select the cluster for which you want to see metrics, or select **All Clusters** to see metrics for all clusters.

**Table**

Each row in the table provides details for an agent.

**Note:** Row columns in this table with an asterisk (\*) at the end of the column definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these columns.

**Agent Count:** The number of agents currently in the table.

**Cluster** The name of the TIBCO BusinessEvents cluster.

**Node** The name of the node.

<b>Agent ID</b>	A unique string that identifies the agent.
<b>Agent Name</b>	The name of the agent.
<b>Type</b>	The type of agent (Inference, Cache, Query, or Dashboard).*
<b>Current State</b>	The current state of the agent.*
<b>Started</b>	When checked, denotes that the agent is started.*
<b>Suspended</b>	When checked, denotes that the agent is suspended.*
<b>Concurrent</b>	When checked, denotes that it is a concurrent agent.*
<b>Queue Capacity</b>	The queue capacity for the agent.*
<b>Queue Size</b>	The queue size for the agent.*
<b>Thread Count</b>	The total number of threads for the agent.*
<b>Total # Rules Fired</b>	The total number of rules fired for the agent.*
<b>Rules/sec</b>	The rate of rules fired for the agent.
<b>Avg Receive Time</b>	See TIBCO documentation for more information.*
<b>Avg Txn Commit Time</b>	The average amount of time taken to commit a transaction.*
<b>Cache Queue Remaining</b>	The total amount of remaining space on the cache queue.*
<b>DB Ops Queue Remaining</b>	The total amount of remaining space on the DB Operations queue.*
<b>Hit Ratio</b>	See TIBCO documentation for more information.*
<b>Job Rate</b>	See TIBCO documentation for more information.*
<b>L1 Cache Max Size</b>	The maximum size of the L1 cache.*
<b>L1 Cache Size</b>	The current size of the L1 cache.*
<b>Max Active</b>	See TIBCO documentation for more information.*
<b># Event Threads</b>	The total number of currently active event threads.*
<b># Jobs</b>	The total number of currently active jobs.*
<b>Priority</b>	See TIBCO documentation for more information.*
<b>Read Only</b>	See TIBCO documentation for more information.*

- Txn Commit Count** The number of transactions committed by the agent.\*
- Txn Receive Count** The number of transactions received by the agent.\*
- Expired** When checked, performance data has not been received within the time specified (in seconds) in the **Expire Time** field in the **Duration** region in the RTView Configuration Application > **(Project Name)** > **Solution Package Configuration** > **TIBCO BusinessEvents** > **DATA STORAGE** tab. The **Delete Time** field (also in the **Duration** region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
- Timestamp** The date and time, relative to the Data Server, that data was last collected for the destination.

## All RTC Reports

Use this display to compare RTC metrics across deployed engines. View a list of all the inference engine RTC reports. You can view reports for a single cluster or all clusters.

The data in this display is identical to the data provided for a single engine in the “[Cluster Summary](#)” display, except that it is aggregated across all inference nodes.

Choose a single cluster or **All Clusters** from the drop-down menus. Each row in the table is a different node.

Cluster	Node	Avg Action	Avg Cache Queue	Avg Cache	Avg DB Ops	Avg DB Qi	Avg DB	Avg Successful
ckfdcache	new51nf	0.00	0.00	0.00	1.00	0.00	8.45	8.33
fdcache	newbe4infern	0.00	0.00	0.00	0.00	0.00	0.00	0.00

**Title Bar** (possible features are):

-   Open the previous and upper display.
-  Open an instance of this display in a new window.
-  Open the online help page for this display.
- open commonly accessed displays.
- The number of items currently in the display.

 **Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

 Open the **Alert Views - RTView Alerts Table** display.

**Filter By:**

The display might include these filtering options:

**Cluster:** Select the cluster for which you want to see metrics, or select **All Clusters** to see metrics for all clusters.

**RTC Txn Manager Reports Table**

Each row in the table is a different report. Data in the row columns describe the report.

**Note:** Row columns in this table with an asterisk (\*) at the end of the column definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these columns.

<b>Report Count:</b>	The number of reports currently in the table.
<b>Cluster</b>	The name of the TIBCO BusinessEvents cluster.
<b>Node</b>	The name of the node.
<b>Avg Action Txn Millisec</b>	The average amount of time taken for an action transaction, in milliseconds.*
<b>Avg Cache Queue Wait Time Millisec</b>	The average cache queue wait time, in milliseconds.*
<b>Avg Cache Txn Millisec</b>	The average amount of time taken for a cache transaction, in milliseconds.*
<b>Avg DB Ops Batch Size</b>	The average database operation batch size.*
<b>Avg DB Queue Wait Time Millisec</b>	The average database queue wait time, in milliseconds.*
<b>Avg DB Txn Millisec</b>	The average amount of time taken for a database transaction, in milliseconds.*
<b>Avg Successful Txn Time Millisec</b>	The average amount of time taken for a successful transaction, in milliseconds.*
<b>Last DB Batch Size</b>	The size of the last database batch.*
<b>Pending Actions</b>	The total number of pending actions.*

<b>Pending Cache Writes</b>	The total number of pending cache writes.*
<b>Pending DB Writes</b>	The total number of pending database writes.*
<b>Pending Events to Ack</b>	The total number of pending events that need to be acknowledged.*
<b>Pending Locks to Release</b>	The total number of pending locks that need to be released.*
<b>Total DB Txns Completed</b>	The total number of database transactions that have been completed.*
<b>Total Successful Txns</b>	The total number of successful transactions.*
<b>Total Errors</b>	The total number of errors.*
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > (Project Name) > <b>Solution Package Configuration</b> > <b>TIBCO BusinessEvents</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Timestamp</b>	The date and time, relative to the Data Server, that data was last collected for the destination.

---

## TIBCO BusinessEvents - HTML

The HTML version features an overview display, "[TIBCO BusinessEvents Overview - HTML](#)" (pictured below), and the following Views which can be found under **Components** tab > **Middleware** > **TIBCO BusinessEvents**:

- ["BE Clusters - HTML"](#)
- ["BE Nodes - HTML"](#)
- ["BE Events - HTML"](#)
- ["BE Concepts - HTML"](#)

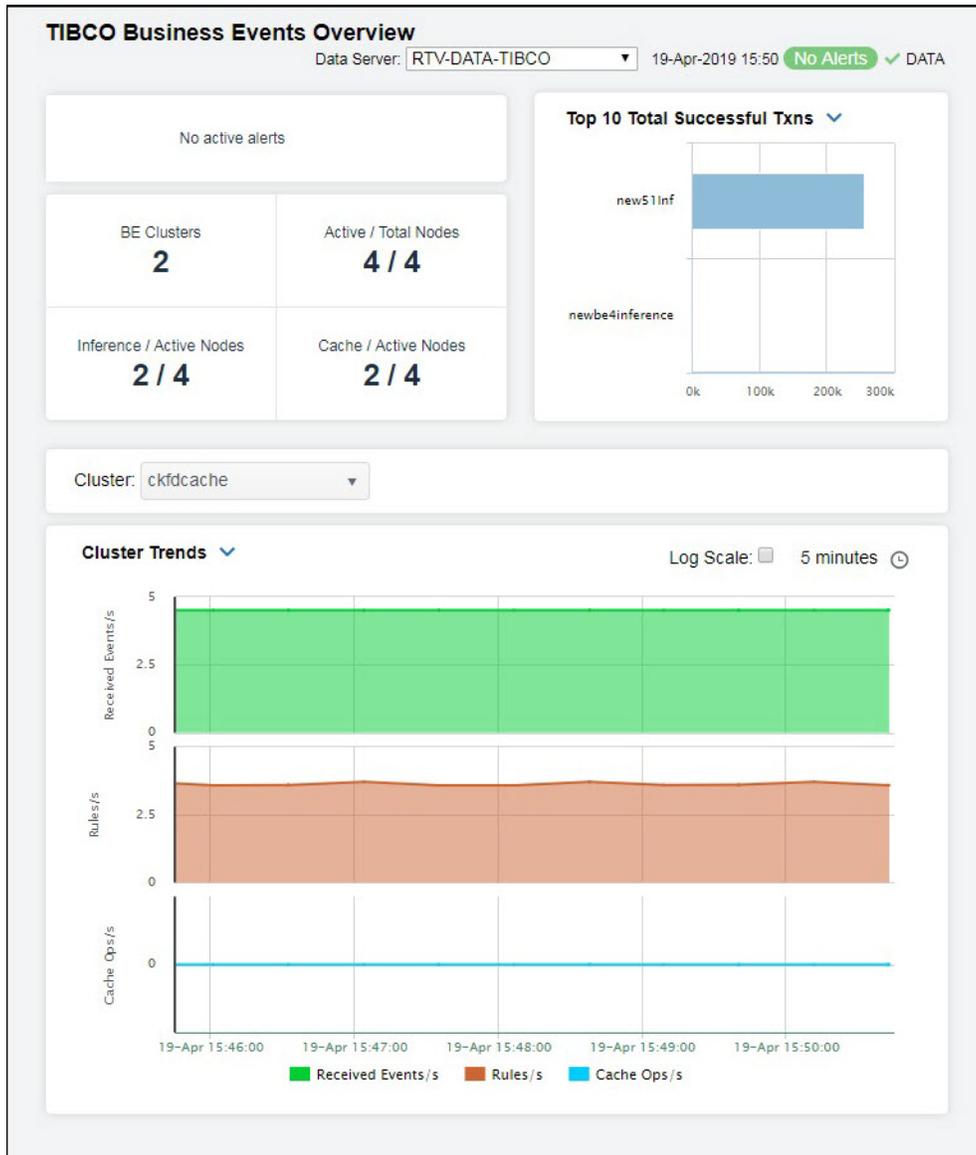
## TIBCO BusinessEvents Overview - HTML

The **TIBCO BusinessEvents Overview** is the top-level display for the TIBCO BusinessEvents Monitor, which provides a good starting point for immediately getting the status of all your clusters, nodes, and transactions on your Data Server. You can select the RTView DataServer for which you want to see data and easily view the current data for that DataServer including:

- The total number of active alerts for the selected DataServer, including the total number of critical and warning alerts.
- The total number of BE clusters.
- The number of active nodes and the total number of nodes.
- The number of inference nodes and cache nodes.
- A visual list of the top 10 servers containing the total successful transactions/total database transactions completed/hit ratio on your connected DataServer.

You can hover over each region in the upper half of the Overview to see more detail. You can also drill down to see even more detail by clicking on each respective region in the Overview. For example, clicking on the alerts in the **CRITICAL** and **WARNING** alerts region opens the **Alerts Table by Components** display.

The bottom half of the display provides Cluster and Backing Store trend graphs for a selected server. You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.



## BE Clusters - HTML

These displays present performance metrics and alert status for your BusinessEvents system. Clicking **BE Clusters** from the left/navigation menu opens the ["TIBCO BE Clusters Table - HTML"](#) display, which shows all available utilization metrics for all BE clusters. The options available under **BE Clusters** are:

- **Clusters Heatmap**: Opens the ["TIBCO BE Clusters Heatmap - HTML"](#), which shows cluster and alert status for all BE clusters.
- **Clusters Summary**: Opens the ["TIBCO BE Cluster Summary - HTML"](#) display, which shows information for a single BE Cluster.

## TIBCO BE Clusters Table - HTML

Use this display to check event, concept, and backing store metrics for all of your clusters. Consider keeping this display open to monitor your TIBCO BusinessEvents clusters in general. Each row in the table contains data for a particular cluster. Click a column header to sort column data in ascending or descending order. Double-click on a table row to drill-down to the ["TIBCO BE Cluster Summary - HTML"](#) display and view metrics for that particular cluster. Toggle between the commonly accessed **Table** and **Heatmap** displays by clicking the drop down list on the display title.

**TIBCO BE Clusters Table** ▾ 22-Apr-2019 10:06 No Alerts ✓ DATA

Clusters: 2

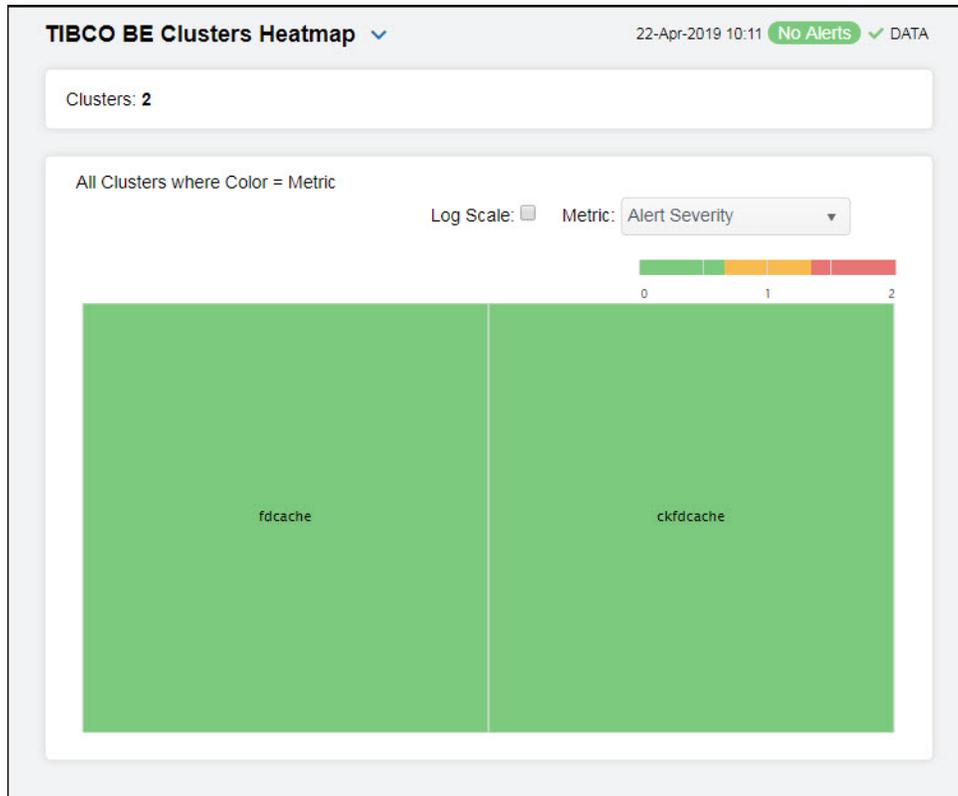
**Clusters**

Cluster Name	Alert Level	Alert Count	Member Count	Events Received ▾	Events Sent	Rec Eve
ckfdcache	✓	0	2	162,650	0	
fdcache	✓	0	2	0	0	

## TIBCO BE Clusters Heatmap - HTML

Clicking **Clusters Heatmap** in the left/navigation menu opens the **TIBCO BE Clusters Heatmap**, which allows you to view the status and alerts of all BE clusters. Use the **Metric** drop-down menu to view the heatmap using a different metric.

The heatmap is organized so that each rectangle represents a cluster. The rectangle color indicates the most critical alert state. Click on a node to drill-down to the "[TIBCO BE Cluster Summary - HTML](#)" display and view metrics for a particular cluster. Toggle between the commonly accessed **Table** and **Heatmap** displays by clicking the drop down list on the display title. Mouse-over rectangles to view more details about cluster performance and status.



### Available Metrics

Select the metric driving the heatmap display. The default is **Alert Severity**. Each **Metric** has a color gradient bar that maps values to colors. The heatmap is organized by clusters, where each rectangle represents a cluster. Mouse-over any rectangle to display the current values of the metrics for the cluster. Click on a rectangle to drill-down to the associated "[TIBCO BE Cluster Summary - HTML](#)" display for a detailed view of metrics for that particular cluster.

#### Alert Severity

The maximum level of alerts in the heatmap rectangle. Values range from **0 - 2**, as indicated in the color gradient bar, where **2** is the highest Alert Severity:

- Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
- Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
- Green indicates that no metrics have exceeded their alert thresholds.

- Alert Count** The total number of critical and warning alerts in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.
- Member Count** The total number of members in the cluster. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum number of members in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
- Events Received** The number of events received in the cluster. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the most events received in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
- Events Sent** The number of events sent in the cluster. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the most events sent in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

## TIBCO BE Cluster Summary - HTML

Clicking **Cluster Summary** in the left/navigation menu opens the **TIBCO BE Cluster Summary** display, which allows you to view configuration and utilization data for a single cluster. Select a cluster to view Rete statistics, cache metrics, Backing Store data, and trend data for the cluster. Clicking on the information boxes at the top of the display takes you to the "TIBCO BE Clusters Table - HTML" display or the "TIBCO BE Events Table - HTML" display, where you can view additional cluster/event data. There are two options in the trend graph region: **Cluster Trends** and **Backing Store Trends**. In the **Cluster Trends** trend graph region, you can view received event rate, rules fired rate, and cache operations rate over a selected time range. In the **Backing Store Trends** trend graph region, you can view backing store operations rate and average maximum time per backstore operation over a selected time range. Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display.



## BE Nodes - HTML

These displays present performance metrics and alert status for your BusinessEvents nodes. Clicking **BE Nodes** from the left/navigation menu opens the ["TIBCO BE Cluster Nodes Table - HTML"](#) display, which shows all available utilization metrics for all BE nodes. The options available under **BE Nodes** are:

- **Cluster Nodes Heatmap:** Opens the ["TIBCO BE Cluster Nodes Heatmap - HTML"](#), which shows cluster and alert status for all BE cluster nodes.
- **Inference Node Summary:** Opens the ["TIBCO BE Inference Node Summary - HTML"](#) display, which shows information for (inference) agents for a single BE cluster node.
- **Storage Node Summary:** Opens the ["TIBCO BE Storage Node Summary - HTML"](#) display, which displays cache data for a specific node.

## TIBCO BE Cluster Nodes Table - HTML

Use this display to view configuration and utilization data for nodes in a cluster. Each row in the table contains data for a particular node. Click a column header to sort column data in ascending or descending order. Toggle between the commonly accessed **Table**, **Heatmap**, and **Summary** displays by clicking the drop down list on the display title.

Cluster Name	Connection	Alert Level	Alert Count	MemberCount	Auto St
ckfdcachecache	new51Cache	OK	0	2	▼
ckfdcachecache	new51Inf	OK	0	2	▼
fdccache	newbe4cachecache	OK	0	2	▼
fdccache	newbe4inference	OK	0	2	▼

## TIBCO BE Cluster Nodes Heatmap - HTML

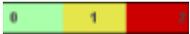
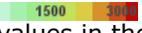
Clicking **Cluster Nodes Heatmap** in the left/navigation menu opens the **TIBCO BE Cluster Nodes Heatmap**, which allows you to view utilization data for all nodes in a cluster in a heatmap format. Use the **Metric** drop-down menu to view the heatmap based on a different metric.

The heatmap is organized so that each rectangle represents a node. The rectangle color indicates the most critical alert state. Toggle between the commonly accessed **Table** and **Heatmap** displays by clicking the drop down list on the display title. Mouse-over rectangles to view more details about cluster performance and status.



### Available Metrics

Select the metric driving the heatmap display. The default is **Alert Severity**. Each **Metric** has a color gradient bar that maps values to colors. The heatmap is organized by nodes, where each rectangle represents a node. Mouse-over any rectangle to display the current values of the metrics for the node.

<b>Alert Severity</b>	<p>The maximum alert level in the item (index) associated with the rectangle. Values range from <b>0</b> to <b>2</b>, as indicated in the color gradient bar , where <b>2</b> is the greatest <b>Alert Severity</b>.</p> <p><b>2</b> -- Metrics that have exceeded their specified <b>ALARMLEVEL</b> threshold and have an Alert Severity value of <b>2</b> are shown in red. For a given rectangle, this indicates that one or more metrics have exceeded their alarm threshold.</p> <p><b>1</b> -- Metrics that have exceeded their specified <b>WARNINGLEVEL</b> threshold and have an Alert Severity value of <b>1</b> are shown in yellow. For a given rectangle, this indicates that one or more metrics have exceeded their warning threshold.</p> <p><b>0</b> -- Metrics that have not exceeded either specified threshold have an Alert Severity value of <b>0</b> and are shown in green. For a given rectangle, this indicates that no metrics have exceeded a specified alert threshold.</p>
<b>Alert Count</b>	<p>The total number of alarm and warning alerts in a given item (index) associated with the rectangle.</p> <p>The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>JVM % CPU Used</b>	<p>The total percentage of JVM CPU used in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from <b>0</b> to the alert threshold of <b>JvmCpuPercentHigh</b>. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>JVM % Memory Used</b>	<p>The total percentage of JVM Memory Used in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from <b>0</b> to the alert threshold of <b>JvmMemoryUsedHigh</b>. The middle value in the gradient bar indicates the middle value of the range.</p>

## TIBCO BE Inference Node Summary - HTML

Clicking **Inference Node Summary** in the left/navigation menu opens the **TIBCO BE Inference Node Summary** display, which allows you to view configuration and utilization data for a single inference node. View a list of all agents on the node and trend graphs tracing the rule execution rate for agents on the node. The rule execution rate is relative to the overall CPU and heap utilization for the engine's JVM.

**NOTE:** An inference node (also known as an engine or processing unit) is the container where one or more inference agents run. Generally, the agents in a given node implement different rule sets, and distributing nodes on different hosts provides fault tolerance and load balancing for the cluster. For details, refer to TIBCO documentation.

Clicking on the information boxes at the top of the display takes you to the ["TIBCO BE Clusters Table - HTML"](#) display, the ["TIBCO BE Agent Event Summary - HTML"](#) display, or the **JVM Summary** display, where you can view additional cluster/event/JVM data. The **Agents for this Node** region lists the inference agents associated with the selected node. There are two options in the trend graph region: **Utilization** and **Rules/s and Threads**. In the **Utilization** trend graph region, you can view percentage of CPU being used by the engine/process and the amount of memory, in megabytes, in use by the JVM for heap space over a selected time range. In the **Rules/s and Threads** trend graph region, you can view the rate of rules fired for the agent and the total number of threads for the agent over a selected time range. Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display.

### TIBCO BE Inference Node Summary 22-Apr-2019 13:38 No Alerts DATA

Node: new51Inf Agent ID: 1

Queue Capacity %  
**0.0**

Top Rules/s  
**3.6**

Top Job Rate  
**4.4**

Top Hit Ratio  
**0.0**

Heap Used %  
**8.9**

CPU %  
**1.8**

#### Agents for this Node

Agent ID	Agent Class	Type	Current State	Started	Suspended	Concurrent
1	inference-class	INFERENCE	ACTIVATED	✓		

#### Utilization Log Scale: 5 minutes ⌚

The utilization charts show CPU % (green area) and Heap usage (blue areas) from 22-Apr 13:34:00 to 22-Apr 13:38:00. The CPU % chart shows a peak of approximately 3.2% at 13:37:30. The Max Heap (light blue) is constant at 1000 MB, and the Used Heap (dark blue) is constant at approximately 100 MB.

Connection: **192.168.200.144:58701** [Critical/Warning: 0/0](#)  
 Expired: **false** Max Number of Jobs: **220,210**  
 Max Number of Rules Fired: **176,168** Max Threads: **10**  
 Max Number of Event Threads: **10** Max Avg Receive Time ms: **0.0** Max Avg Commit Time ms: **8.0**

Last Update: **22-Apr-2019 13:38:36**

## TIBCO BE Storage Node Summary - HTML

Clicking **Storage Node Summary** in the left/navigation menu opens the **TIBCO BE Storage Node Summary** display, which allows you to view configuration details for a single cache node and a list of all caches that are backed by the backing store (database).

**NOTE:** A storage node (also known as a cache node) provides fast access to events and concepts required during each RTC by the inference engines. Storage nodes also serve as buffers for reads and writes between the cluster and the backing store. For details, refer to TIBCO documentation.

Clicking on the information boxes at the top of the display takes you to the **JVM Summary** display, where you can view additional JVM data. The **Backing Store** region lists the caches that are backed by the backing store in the selected node. There are two options in the trend graph region: **Utilization** and **Table Sizes**. In the **Utilization** trend graph region, you can view percentage of CPU being used by the node and the amount of memory, in megabytes, in use by the JVM for heap space over a selected time range. In the **Table Sizes** trend graph region, you can view the number of unique objects cached in the local index table and the number of entries in the table of external IDs used as indexes by the backing store over a selected time range. Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display.

### TIBCO BE Storage Node Summary 17-May-2019 15:34 No Alerts DATA

Node: new51Cache

CPU %  
**1.1**

Top Load Avg Time ms  
**0.0**

Top Store Avg Time ms  
**0.0**

Heap Used %  
**4.8**

Top Loads/s  
**0.0**

Top Stores/s  
**0.0**

#### Backing Store

Cache Name	Active	Delete Avg Time	Load Avg Time	Store Avg Time	Delete Total
be_gen_Events_Deposit	✓	0.0	0.0	0.0	0
be_gen_Concepts_Account	✓	0.0	0.0	0.0	0
be_gen_Events_Debit	✓	0.0	0.0	0.0	0
be_gen_Events_Unsuspend	✓	0.0	0.0	0.0	0
be_gen_FraudCriteria	✓	0.0	0.0	0.0	0
com_tibco_cep_runtime_model_element	✓	0.0	0.0	0.0	0

#### Utilization Log Scale: 15 minutes ⌵

■ CPU %    ■ Max Heap    ■ Used Heap

Connection: **192.168.200.144:58700**    Critical/Warning: 0/0

Expired: **false**                      Node ID: **7f000001-e4f2**                      Version: **v5.x**

Storage Enabled: **true**

---

Last Update: **17-May-2019 15:34:44**

## BE Events - HTML

These displays present performance metrics and alert status for your BusinessEvents events. Clicking **BE Events** from the left/navigation menu opens the ["TIBCO BE Events Table - HTML"](#) display, which shows all available utilization metrics for all BE events. The options available under **BE Events** are:

- **Event Summary:** Opens the ["TIBCO BE Agent Event Summary - HTML"](#) display, which shows information for a single BE event.
- **Event Cache Hits:** Opens the ["TIBCO BE Event Cache Hits Table - HTML"](#) display, which
- **Event Hit Summary:** Opens the ["TIBCO BE Event Hit Summary - HTML"](#) display, which

### TIBCO BE Events Table - HTML

View run-time statistics for a selected group of agents. With TIBCO BusinessEvents, events are cached when they are out of scope, and deleted or persisted to the backing store when they are no longer useful. Double-clicking on a row in the table displays access patterns over time for the event in the ["TIBCO BE Agent Event Summary - HTML"](#) display. Each row in the table contains data for a particular event. Click a column header to sort column data in ascending or descending order. Toggle between the commonly accessed **Table** and **Summary** displays by clicking the drop down list on the display title.

**NOTE:** Events cause rules to execute in the BusinessEvents Rete network. Events can be created by external phenomena, such as the arrival of a JMS message, or internally when rules are processed. When an event enters the Rete network, it causes a run-to-completion cycle which continues until no further rules can be processed. Each named event that can be handled by a BusinessEvents application is specified at build time in BusinessEvents studio. For details, refer to TIBCO documentation.

**TIBCO BE Events Table** 22-Apr-2019 14:23 DATA

Node: - All - Agent ID: - All -

Events: 9

**Agent Events**

Event	Node	Avg Time Pre RTC s	Avg Time in RTC s
AccountOperations	newbe4inference	0.0	
Unsuspend	newbe4inference	0.0	
Debit	newbe4inference	0.0	
CreateAccount	newbe4inference	0.0	
Deposit	new51Inf	0.99	26
AccountOperations	new51Inf	0.0	
Unsuspend	new51Inf	0.79	21
Debit	new51Inf	1.55	25
CreateAccount	new51Inf	0.0	

## TIBCO BE Agent Event Summary - HTML

Clicking **Event Summary** in the left/navigation menu opens the **TIBCO BE Agent Event Summary** display, which allows you to view detailed performance metrics for an agent's event.

Clicking on the information boxes at the top of the display takes you to the "[TIBCO BE Events Table - HTML](#)" display, where you can view additional event data. There are three options in the trend graph region: **Time in RTC**, **Agent Rates**, and **Channel Rates**. In the **Time in RTC** trend graph region, you can view the average amount of time taken for the event to begin its run to completion cycle, the average amount of time taken for the event to complete (once it has started) its run to completion cycle, and the average amount of time taken by the event after its run to completion cycle has ended over a selected time range. In the **Agent Rates** trend graph region, you can view rate of event assertions into the Rete network via the agent, the rate of event retractions/deletions from the Rete network via the agent, and the rate of event modifications in the Rete network via the agent over a selected time range. In the **Channel Rates** trend graph region, you can view rate of event assertions into the Rete network via the channel, the rate of event retractions/deletions from the Rete network via the channel, and the rate of event modifications in the Rete network via the channel over a selected time range. Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display.

TIBCO BE Agent Event Summary ▼
24-May-2019 15:39 No Alerts ✓ DATA ↗ ?

Node: new51Inf ▼
Agent ID: 1 ▼

Event: Unsuspend ▼

Assertions/s Agents / Channel  
**0.0 / 0.9**

Retractions/s Agents / Channel  
**0.0 / 0.9**

Modified/s Agents / Channel  
**0.0 / 0.0**

Avg Time in RTC s  
**41.9**

Avg Time Pre / Post RTC s  
**0.7 / 0.0**

Hits / Misses in L1 Cache  
**0 / 0**

Time in RTC ▼
Log Scale:  15 minutes 🕒

■ Avg Time Pre RTC s
■ Avg Time In RTC s
■ Avg Time Post RTC s

Last Update: 24-May-2019 15:40:10 [Critical/Warning: 0/0](#)

Cache Mode: CACHE      Current Events Asserted From Channel: 27

Current Events Retracted From Channel: 27      Current Events Modified From Channel: 0

Current Events Asserted From Agents: 0      Current Events Retracted From Agents: 0

Current Events Modified From Agents: 0

## TIBCO BE Event Cache Hits Table - HTML

Clicking **Event Cache Hits** in the left/navigation menu opens the **TIBCO BE Event Cache Hits Table** display, which allows you to view cache performance metrics per event for a single cluster or **All Clusters**.

**NOTE:** Events cause rules to execute in the BusinessEvents Rete network. Events can be created by external phenomena, such as the arrival of a JMS message, or internally when rules are processed. When an event enters the Rete network, it causes a run-to-completion cycle which continues until no further rules can be processed. Each named event that can be handled by a BusinessEvents application is specified at build time in BusinessEvents studio. For details, refer to TIBCO documentation.

Click a column header to sort column data in ascending or descending order. Double-click on a table row to drill-down to the ["TIBCO BE Event Hit Summary - HTML"](#) display and view metrics for that particular event. Toggle between the commonly accessed **Table** and **Summary** displays by clicking the drop down list on the display title.

**TIBCO BE Event Cache Hits Table** ▾
22-Apr-2019 15:02 No Alerts ✓ DATA

Cluster: - All -
Nodes: new51Inf

Events: **5**

**Node-Level Event Cache Hits**

Event	Node	Cache Size	Get Avg Time	Put Avg Time	Re Av
AccountOperations	new51Inf	0	0.0	0.0	
CreateAccount	new51Inf	0	0.0	0.0	
Debit	new51Inf	0	0.0	0.0	
Deposit	new51Inf	0	0.0	0.0	
Unsuspend	new51Inf	0	0.0	0.0	

## TIBCO BE Event Hit Summary - HTML

Clicking **Event Hit Summary** in the left/navigation menu opens the **TIBCO BE Event Hit Summary** display, which allows you to view detailed event performance metrics for a single cluster or **All Clusters**, a node, and an event.

**NOTE:** Events cause rules to execute in the BusinessEvents Rete network. Events can be created by external phenomena, such as the arrival of a JMS message, or internally when rules are processed. When an event enters the Rete network, it causes a run-to-completion cycle which continues until no further rules can be processed. Each named event that can be handled by a BusinessEvents application is specified at build time in BusinessEvents studio. For details, refer to TIBCO documentation.

There are two options in the trend graph region: **Hit Rates** and **Current Hits**. In the **Hit Rates** trend graph region, you can view the number of "get" operations per second, the number of "put" operations per second, and the number of "remove" operations per second over a selected time range. In the **Current Hits** trend graph region, you can view the number of "get" operations, the number of "put" operations, and the number of "remove" operations over a selected time range. Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display.



## BE Concepts - HTML

These displays present performance metrics and alert status for your BusinessEvents concepts. Clicking **BE Concepts** from the left/navigation menu opens the ["TIBCO BE Concepts Table - HTML"](#) display, which shows all available utilization metrics for all BE concepts. The options available under **BE Concepts** are:

- **Concept Hit Summary:** Opens the ["TIBCO BE Concept Hit Summary - HTML"](#), which shows details and alert status for a BE concept.
- **Channels:** Opens the ["TIBCO BE Channel Status Table - HTML"](#) display, which shows information for destinations and channels for a single BE node.
- **Inference Agents Table:** Opens the ["TIBCO BE Inference Agents Table - HTML"](#) display, which displays agents data for a specific cluster.
- **RTC Reports Table:** Opens the ["TIBCO BE RTC Txn Manager Reports - HTML"](#) display, which displays reports data for a specific cluster.

## TIBCO BE Concepts Table - HTML

View a list of concepts and their run-time statistics. Choose a single cluster or **All Clusters** and a node from the drop-down menus. Double-clicking on a row in the table displays additional details as well as hit rates and current hits over time for the concept in the ["TIBCO BE Concept Hit Summary - HTML"](#) display. Each row in the table contains data for a particular concept. Click a column header to sort column data in ascending or descending order. Toggle between the commonly accessed **Table** and **Summary** displays by clicking the drop down list on the display title.

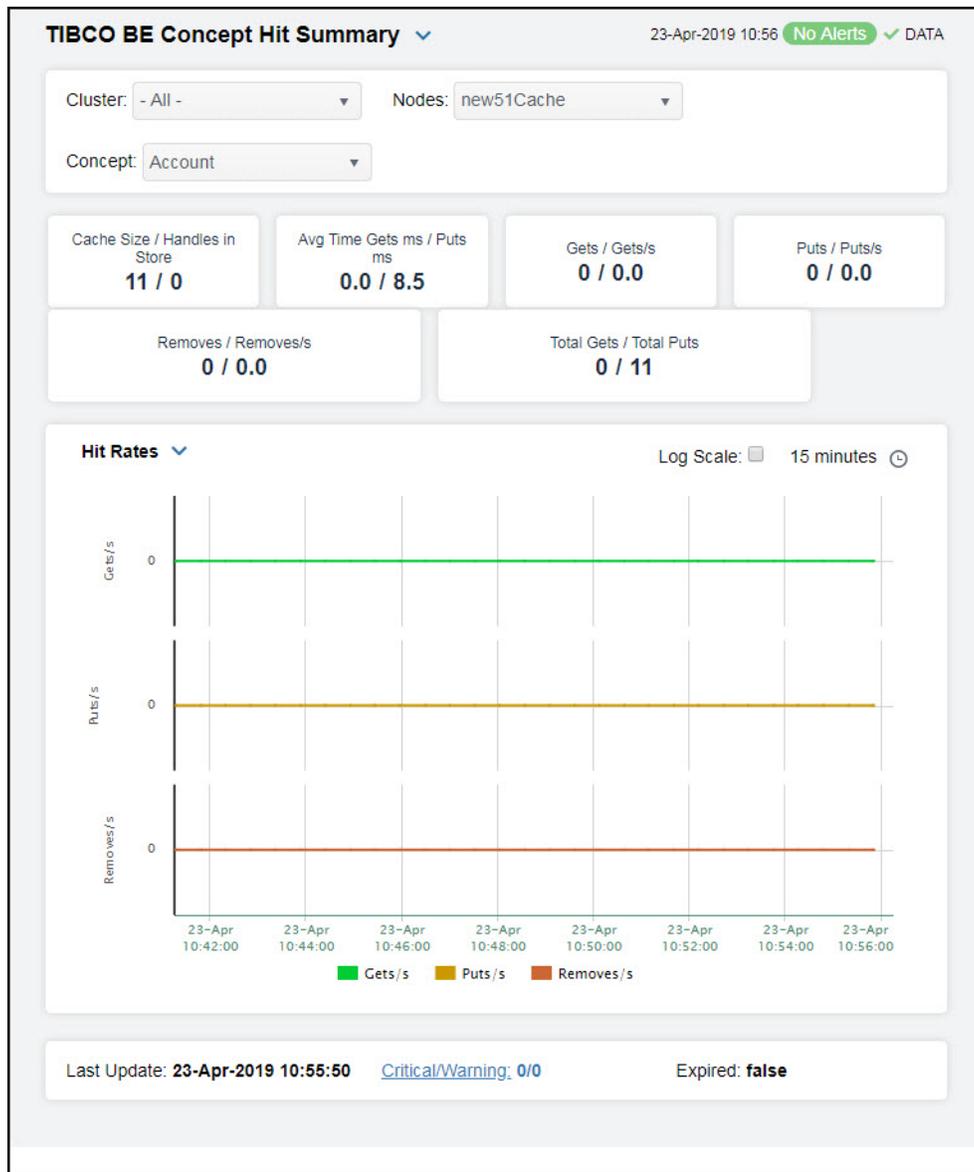
Concept	Connection	Cache Size	Get Avg Time	Put Avg Time	Ren Avg
Account	new51Cache	11	0.0	8.45	

## TIBCO BE Concept Hit Summary - HTML

Clicking **Concept Hit Summary** in the left/navigation menu opens the **TIBCO BE Concept Hit Summary** display, which allows you to view current and historic data for a single concept. Data in this display can be useful if your BusinessEvents system uses Cache object management. When Cache object management is used, concepts with a sufficiently long time to live (TTL) setting are cached.

Cache reference patterns for certain concepts may be related to incoming events (for example, customer purchase orders with associated inventory queries). The trend charts show the cache activity of such concepts, and might be useful in diagnosing the behavior of your application over time.

There are two options in the trend graph region: **Hit Rates** and **Current Hits**. In the **Hit Rates** trend graph region, you can view the number of "get" operations per second, the number of "put" operations per second, and the number of "remove" operations per second over a selected time range. In the **Current Hits** trend graph region, you can view the number of "get" operations, the number of "put" operations, and the number of "remove" operations over a selected time range. Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display.



## TIBCO BE Channel Status Table - HTML

Clicking **Channels** in the left/navigation menu opens the **TIBCO BE Channel Status Table** display, which allows you to view a list of destinations, which are sources and sinks of events, and Channels. Destinations are potentially bi-directional, and the table indicates whether events are sent or received. Channels provide a class wrapper for destinations, and make it possible to enable or disable a group of destinations with one operation.

Click a column header to sort column data in ascending or descending order. Toggle between the commonly accessed **Table** and **Summary** displays by clicking the drop down list on the display title.

**TIBCO BE Channel Status Table** 23-Apr-2019 10:58 ✓ DATA

Cluster: - All - Nodes: new51Inf

Destinations: 1 Channels: 1

**Destinations**

Destination URI	Suspended	Num Events Received	Num Events Sent	Rece Even
/Channels/HTTP/AllOps		177,106	0	

**Channels**

Channel URI	State	Expired
/Channels/HTTP	Started	

## TIBCO BE Inference Agents Table - HTML

Clicking **Inference Agents Table** in the left/navigation menu opens the **TIBCO BE Inference Agents Table** display, which allows you to compare agent metrics across deployed engines and verify that the cluster is properly load-balanced. View a list of all the inference agents deployed in each cluster. You can view agent data for a single cluster or all clusters.

The data in this display is identical to the data provided for a single engine in the ["TIBCO BE Cluster Summary - HTML"](#) display, except that it is aggregated across all inference nodes.

Choose a single cluster or **All** clusters from the drop-down menu. Each row in the table is a different agent.

**TIBCO BE Inference Agents Table** 17-May-2019 15:40 ✓ DATA

Cluster: - All -

Agents: 2

**Inference Agents**

Cluster Name	Node	Agent ID	Agent Class	Type
ckfdcache	new51Inf	0	inference-class	INFERENCE
fdcache	newbe4inference	2	inference-class	Inference

## TIBCO BE RTC Txn Manager Reports - HTML

Clicking **RTC Reports Table** in the left/navigation menu opens the **TIBCO BE RTC Txn Manager Reports** display, which allows you to compare RTC metrics across deployed engines. View a list of all the inference engine RTC reports. You can view reports for a single cluster or all clusters. The data in this display is identical to the data provided for a single engine in the ["TIBCO BE Cluster Summary - HTML"](#) display, except that it is aggregated across all inference nodes.

**TIBCO BE RTC Txn Manager Reports** 23-Apr-2019 11:24 ✓ DATA

Cluster: - All -

Reports: 2

**RTC Txn Manager Reports**

Cluster Name	Node	Avg Action Txn ms	Avg Cache Queue Wait Time ms	Avg Cache Txn ms
ckfdcache	new51Inf	0.0	0.0	
fdcache	newbe4inference	0.0	0.0	

## TIBCO BusinessWorks

The following TIBCO BusinessWorks Views can be found under **Components** tab > **Middleware**. The displays within the Views will be populated with data once the Solution Package for TIBCO BusinessWorks is configured in the RTView DataServer for TIBCO and the RTView DataServer for TIBCO is connected to RTViewCentral.

- ["BW Applications"](#)
- ["BW Containers"](#)
- ["BW AppNodes"](#)
- ["BW AppSlices"](#)
- ["BW Processes"](#)
- ["BW5 Engines"](#)
- ["BW5 Processes"](#)
- ["BW5 Activities"](#)
- ["BW5 Servers"](#)

## BW Applications

These displays present process performance data for your BusinessWorks applications and AppSpaces across BusinessWorks Domains. Process metrics are totaled by application. Use these displays to monitor critical alerts for all your BusinessWorks applications, and investigate those alerts in lower-level displays. Displays in this View are:

- ["BW All Applications Heatmap" on page 1248](#): A color-coded heatmap view of selected application performance metrics.
- ["BW All Applications Table" on page 1251](#): A tabular view of all available application performance data in this BusinessWorks View.
- ["BW Single Application Summary" on page 1254](#): Current and historical metrics for a single application.

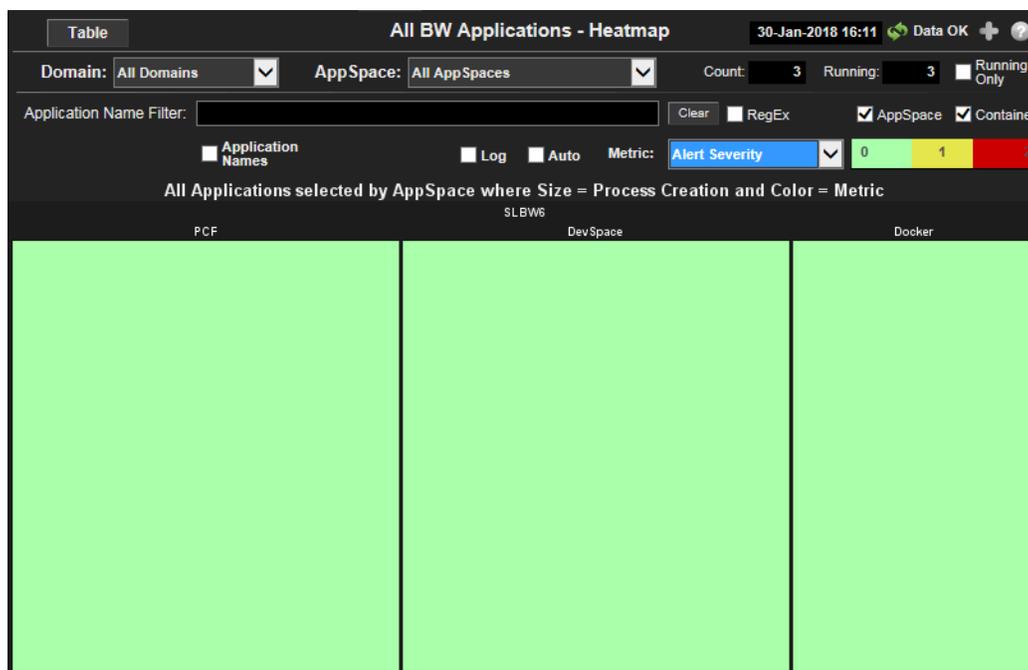
### BW All Applications Heatmap

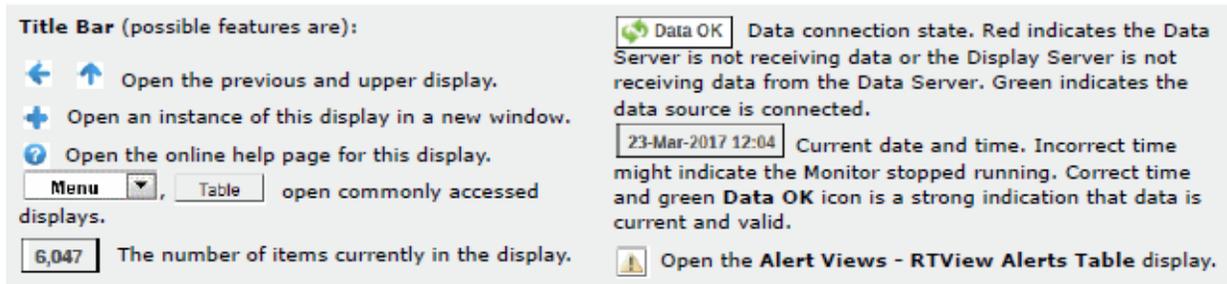
View the most critical BusinessWorks application alert states pertaining to process creation and execution for all nodes on which the applications are deployed. Use this display to quickly identify applications with critical alerts.

Each rectangle in the heatmap represents an application. The rectangle color indicates the most critical alert state associated with the application. The rectangle size represents process creation across applications; a larger size is a larger value.

Choose a domain and AppSpace from the drop-down menus. Choose a different metric to display from the **Metric** drop-down menu. Enter a string in the **Application Name Filter** field to limit data shown in the display. Use the **Application Names** check-box  to include or exclude labels in the heatmap. Mouse over a rectangle to see additional metrics. By default, this display shows **Alert Severity**.

Drill-down and investigate an application by clicking a rectangle in the heatmap to view details in the ["BW Single Application Summary"](#) display.



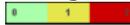
**Filter By:**

The display might include these filtering options:

- Domain:** Select the domain for which you want to view data in the display.
- AppSpace** Select the AppSpace for which you want to view data in the display.
- Application Name Filter** Enter a string (all or part of a application name) to filter the data shown in the display. If you enter part of an application name, you must enter "\*" before and/or after the string. For example, if you have an application named AppNameOne, you could filter using \*Name\*, \*NameOne, or AppName\*.
- Clear** Clears the **Application Name Filter** entries from the display.
- RegEx** Toggles the **Application Name Filter** to accept Regular Expressions for filtering. For example, if your application name is AppNameOne and this option was toggled on, you could enter "Name" (without using "\*" to display the application in the heatmap).
- AppSpace** When selected, those AppNodes deployed in an AppSpace display in the heatmap.
- Container** When selected, those AppNodes deployed in a container display in the heatmap.
- Application Names** Check to include labels in the heatmap.

**Fields and Data:**

- Count:** The total number of AppSpaces currently shown in the display.
- Running** The total number of AppSpaces currently running in the display.
- Running Only** Select to show only running applications in the display.
- Log** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Auto** Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value. NOTE: Some metrics auto-scale automatically, even when **Auto** is not selected.
- Metric** Choose a metric to view in the display.

<b>Alert Severity</b>	<p>The maximum level of alerts in the heatmap rectangle. Values range from <b>0</b> - <b>2</b>, as indicated in the color gradient  bar, where <b>2</b> is the highest Alert Severity:</p> <ul style="list-style-type: none"> <li> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Alert Count</b>	<p>The total number of critical and warning alerts in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.</p>
<b>Active Count</b>	<p>The total number of active processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.</p>
<b>Completed Count</b>	<p>The total number of completed processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.</p>
<b>Suspended Count</b>	<p>The total number of suspended processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.</p>
<b>Failed Count</b>	<p>The total number of failed processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.</p>
<b>Created / sec</b>	<p>The number of processes created per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.</p>
<b>Suspended / sec</b>	<p>The number of suspended processes per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.</p>
<b>Failed / sec</b>	<p>The number of failed processes per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.</p>

<b>Exec Time / sec</b>	The process execution time per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Most Recent Exec Time</b>	The execution time for the most recently executed process in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Average Exec Time</b>	The average execution time for all processes in the heatmap rectangle, calculated by dividing the delta execution time for the interval by the delta completed, or the number of process instances that completed in the interval. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Average Elapsed Time</b>	The average elapsed time for all processes in the heatmap rectangle, calculated by dividing the delta elapsed time for the interval by the delta completed, or the number of process instances that completed in the interval. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

## BW All Applications Table

View BusinessWorks data shown in the ["BW All Applications Heatmap"](#), and additional details, in a tabular format.

Each row in the table is an application. Choose a domain and AppSpace from the drop-down menus. Click a column header to sort column data in numerical or alphabetical order. Enter a string in the **Application Name Filter** field to limit data shown in the display.

Drill-down and investigate by clicking a row to view details for the selected application in the ["BW Single Application Summary"](#) display.

Domain	AppSpace	Name	Alert Level	Alert Count	State	Deployment	AppSource
SLBW6	DevSpace	tibco.bw.sample.binding.rest.BookStore.applicati	Green		Running	Appspace	
SLBW6	Docker	tibco.bwce.sample.binding.rest.BookStore.applic:	Green		Running	Container	
SLBW6	PCF	tibco.bwce.sample.binding.rest.BookStore.applic:	Green		Running	Container	
standalone	standalone-34	tibco.bwce.sample.binding.rest.BookStore.applic:	Green		Running	Container	
standalone	standalone-9f	tibco.bwce.sample.binding.rest.BookStore.applic:	Green		Running	Container	
standalone	standalone-c5	tibco.bwce.sample.binding.rest.BookStore.applic:	Green		Running	Container	

#### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** **Table** open commonly accessed displays.
- 6,047** The number of items currently in the display.



**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04**

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.



Open the **Alert Views - RTView Alerts Table** display.

#### Filter By:

The display might include these filtering options:

- Domain:** Choose a domain to show data for in the display.
- AppSpace** Choose an AppSpace to show data for in the display.
- Application Name Filter** Enter a string (all or part of a application name) to filter the data shown in the display. If you enter part of an application name, you must enter "\*" before and/or after the string. For example, if you have an application named AppNameOne, you could filter using \*Name\*, \*NameOne, or AppName\*.
- Clear** Clears the **Application Name Filter** entries from the display.
- RegEx** Toggles the **Application Name Filter** to accept Regular Expressions for filtering. For example, if your application name is AppNameOne and this option was toggled on, you could enter "Name" (without using "\*" to display the application in the table).
- AppSpace** When selected, those AppNodes deployed in an AppSpace display in the table.
- Container** When selected, those AppNodes deployed in a container display in the table.

**Fields and Data:**

<b>Count:</b>	The total number of applications in the AppSpace.
<b>Running</b>	The total number of applications currently running in the AppSpace.
<b>Running Only</b>	Select to show only running applications in the display.

**Table:**

Each row in the table is a different application.

<b>Domain</b>	The domain in which the application resides.
<b>AppSpace</b>	The AppSpace in which the application resides.
<b>Name</b>	The name of the application.
<b>Alert Level</b>	The most critical alert state for alerts in the row:  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
<b>Alert Count</b>	The total number of active alerts for the application.
<b>State</b>	The current status of the application. Valid values are <b>Running</b> and <b>Stopped</b> .
<b>AppNodes</b>	The total number of AppNodes associated with the application.
<b>Active Processes</b>	The number of currently active application processes.
<b>Suspended Processes</b>	The number of suspended application processes.
<b>Failed Processes</b>	The number of failed application processes.
<b>Completed Processes</b>	The number of completed application processes.
<b>Created /sec</b>	The number of application processes created per second.
<b>Suspended / sec</b>	The number of application process suspensions per second.
<b>Failed /sec</b>	The number of application process failures per second.
<b>Exec Time/ sec</b>	The number of processes executed per second.
<b>Recent Exec Time</b>	The number of seconds for the most recently executed process.
<b>Average Exec Time</b>	The average number of seconds for all processes to execute.
<b>Version</b>	The application version.
<b>Module</b>	The application module.
<b>Shared Module</b>	The shared module, if any.
<b>Time Stamp</b>	The date and time the row data was last updated.

- Source** Name of RTView Data Server sending this data (or localhost).
- Expired** When checked, data has not been received from this host in the specified amount of time.

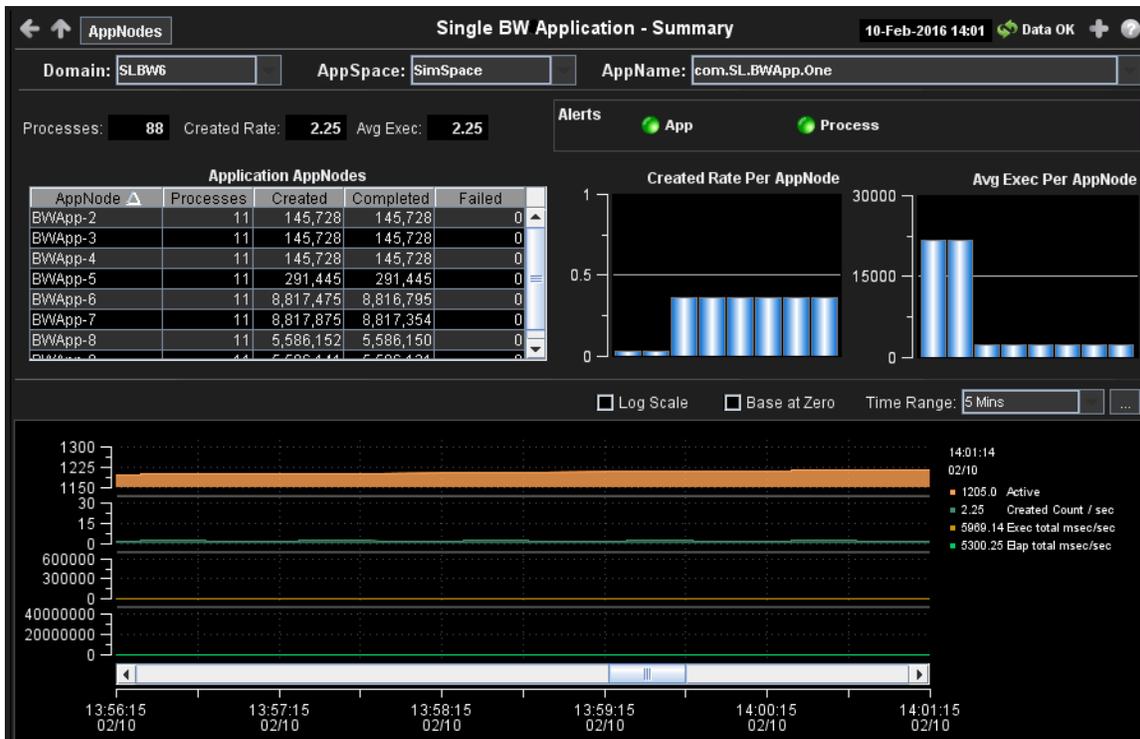
### BW Single Application Summary

View current and historical metrics for a single BusinessWorks application across multiple nodes. Use this display to investigate performance issues of application AppNodes within an AppSpace. Use this display to view all available data for each AppNode by Domain and AppSpace.

This display includes a list of AppNodes with their host names and memory metrics, bar graphs per AppNode for process creation and execution, and trend graphs of process creation and execution metrics.

The summary display also shows the AppNodes of the deployment and process metrics totaled by AppNode. This is useful to see the deployment and load balancing of the Application in current and historical time.

Choose a domain, AppSpace and Application from the drop-down menus. Drill-down and investigate by clicking an AppNode in the table to view details in the "BW Single AppNode Summary" display.



**Title Bar** (possible features are):

- ← ↑ Open the previous and upper display.
- ⊕ Open an instance of this display in a new window.
- ⓘ Open the online help page for this display.
- Menu ▾, Table open commonly accessed displays.
- 6,047 The number of items currently in the display.

 Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

 Open the Alert Views - RTView Alerts Table display.

**Filter By:**

The display might include these filtering options:

- Domain:** Select the domain for which you want to view data in the display.
- AppSpace:** Choose the AppSpace for which you want to view data in the display.
- AppName:** Choose the AppName for which you want to view data in the display.

**Fields and Data:**

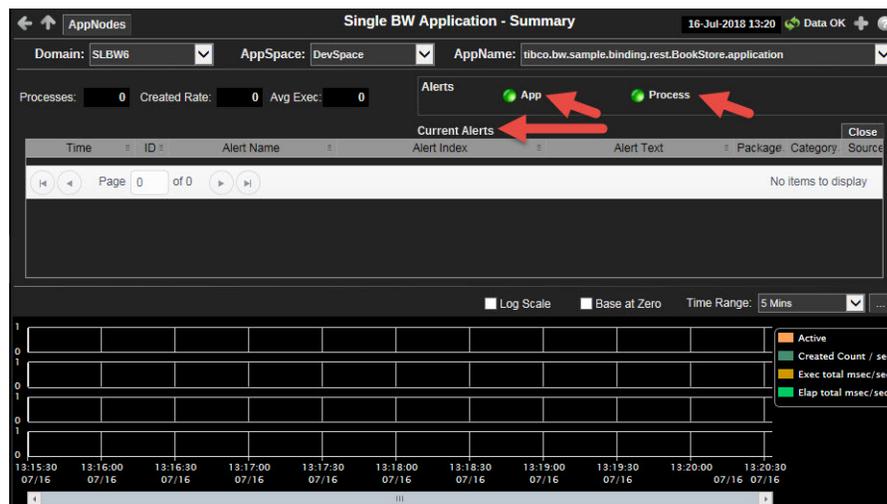
- Processes:** The number of processes currently running for the selected application.
- Created Rate:** The number of processes created per second for the selected application.
- Avg Exec:** The average number of seconds for processes to execute for the selected application.

**Alerts**

Indicates the greatest severity level and the number of open **App** and **Process** alerts for the selected application. Values range from **0** to **2**, where **2** is the greatest Severity:

-  One or more alerts exceeded their ALARM LEVEL threshold.
-  One or more alerts exceeded their WARNING LEVEL threshold.
-  No alert thresholds have been exceeded.

Click on the alert indicator to display a table listing the current alerts for the selected application. Click the **Close** button (for the current alerts table) to close the table.



**Application Appnodes Table**

Each row in the table is a different AppNode. Column values describe processes for the selected application on that AppNode. Click a row to view AppNode details in the ["BW Single AppNode Summary"](#) display.

<b>AppNode</b>	The name of the AppNode.
<b>Processes</b>	The number of processes currently running on the AppNode.
<b>Created</b>	The total number of processes created on the AppNode.
<b>Completed</b>	The total number of completed processes on the AppNode.
<b>Failed</b>	The total number of failed processes on the AppNode.

**Created Rate Per AppNode Bar Graph**

The bar graph shows the current process creation rate per AppNode. Click to drill-down and investigate in the ["BW Single AppSlice Summary"](#) display.

**AvgExec Per AppNode Bar Graph**

The bar graph shows the current average process execution rate per AppNode for the selected application. Click to drill-down and investigate in the ["BW Single AppSlice Summary"](#) display.

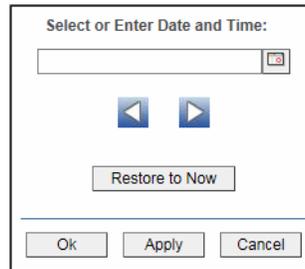
**Trend Graphs**

Traces the sum of process metrics across all processes in all slices of the selected application.

<b>Active</b>	Traces the number of currently active application processes.
<b>Created Count / sec</b>	Traces the number of created application processes.
<b>Exec total msec/sec</b>	Traces the rate at which the application is accumulating process execution time, in milliseconds per second.
<b>Elap total msec/sec</b>	Traces the rate at which the application accumulates process elapsed time, in milliseconds per second.
<b>Log Scale</b>	Select to enable a logarithmic scale. Use <b>Log Scale</b> to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. <b>Log Scale</b> makes data on both scales visible by applying logarithmic values rather than actual values to the data.

**Base at Zero** Select to use zero (0) as the Y axis minimum for all graph traces.

**Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## BW Containers

These displays present process performance data for your BusinessWorks containers across BusinessWorks Domains. Process metrics are totaled by container. Use these displays to monitor critical alerts for all your BusinessWorks containers, and investigate those alerts in lower-level displays. Displays in this View are:

- ["All Containers Heatmap" on page 1257](#): A color-coded heatmap view of selected container performance metrics.
- ["All Containers Table" on page 1260](#): A tabular view of all available container performance data in this BusinessWorks View.
- ["Single Container Summary" on page 1263](#): Current and historical metrics for a single container.

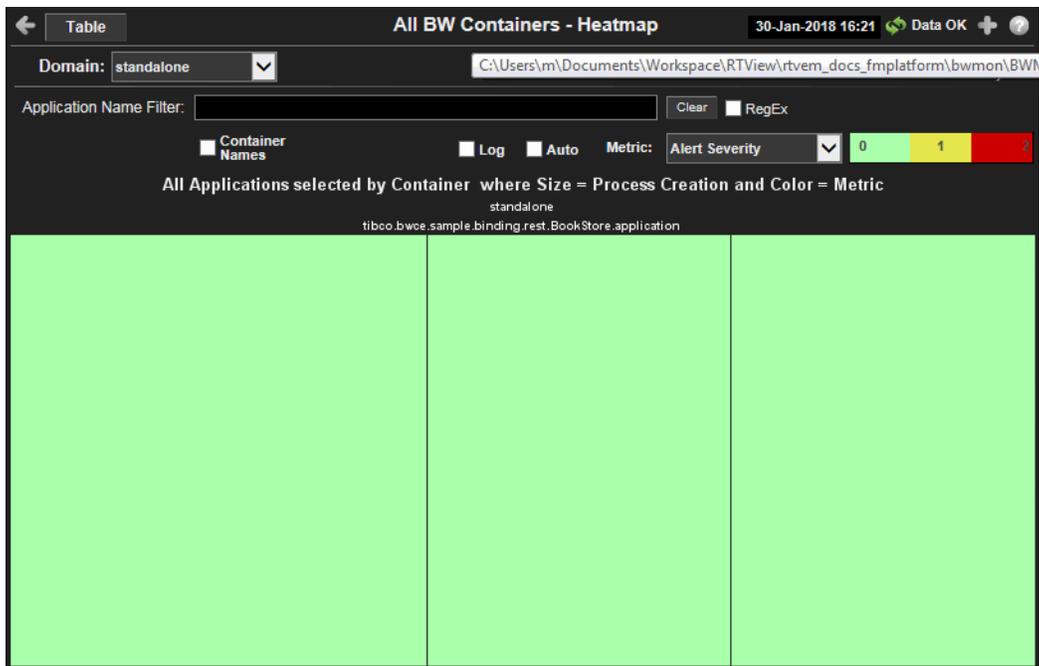
### All Containers Heatmap

View the most critical BusinessWorks container alert states pertaining to process creation and execution for all nodes on which the containers are deployed. Use this display to quickly identify containers with critical alerts.

Each rectangle in the heatmap represents a container. The rectangle color indicates the most critical alert state associated with the container. The rectangle size represents process creation across containers; a larger size is a larger value.

Choose a domain and AppSpace from the drop-down menus containing the containers for which you want to view metrics. By default, this display shows the heatmap based on the **Alert Severity** metric, but you can select a different metric from the **Metric** drop-down menu to view the heatmap based on the selected metric. To view data shown for a specific container(s) in the display, enter a string in the **Container Name Filter** field. Use the **Container Names** check-box  to include or exclude labels in the heatmap. You can mouse over a rectangle to see additional metrics.

Drill-down and investigate a container by clicking a rectangle in the heatmap, which opens the details for the selected container in the **"Single Container Summary"** display.



#### Title Bar (possible features are):

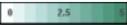
- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.

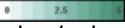
- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

#### Filter By:

The display might include these filtering options:

- Domain:** Select the domain for which you want to view data in the display.
- Application Name Filter** Enter a string (all or part of an application name) to filter the data shown in the display. If you enter part of a application name, you must enter "\*" before and/or after the string. For example, if you have an application named AppNameOne, you could filter using \*Name\*, \*NameOne, or AppName\*. You can also enable the **RegEx** toggle to just enter a portion of the application name.

	<b>Clear</b>	Clears the <b>Application Name Filter</b> entries from the display.
<b>RegEx</b>		Toggles the <b>Application Name Filter</b> to accept Regular Expressions for filtering. For example, if your application name is AppNameOne and this option was toggled on, you could enter "Name" (without using "*" to display the application in the heatmap).
<b>Container Names</b>		Check to include container name labels in the heatmap.
<b>Fields and Data:</b>		
<b>Count:</b>		The total number of containers currently shown in the display.
<b>Running</b>		The total number of containers currently running in the display.
<b>Running Only</b>		Select to show only running containers in the display.
<b>Log</b>		Select to enable a logarithmic scale. Use <b>Log Scale</b> to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. <b>Log Scale</b> makes data on both scales visible by applying logarithmic values rather than actual values to the data.
<b>Auto</b>		Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value. NOTE: Some metrics auto-scale automatically, even when <b>Auto</b> is not selected.
<b>Metric</b>		Choose a metric to view in the display.
	<b>Alert Severity</b>	<p>The maximum level of alerts in the heatmap rectangle. Values range from <b>0</b> - <b>2</b>, as indicated in the color gradient  bar, where <b>2</b> is the highest Alert Severity:</p> <ul style="list-style-type: none"> <li> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
	<b>Alert Count</b>	The total number of critical and warning alerts in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.
	<b>Active Count</b>	The total number of active processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
	<b>Completed Count</b>	The total number of completed processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
	<b>Suspended Count</b>	The total number of suspended processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

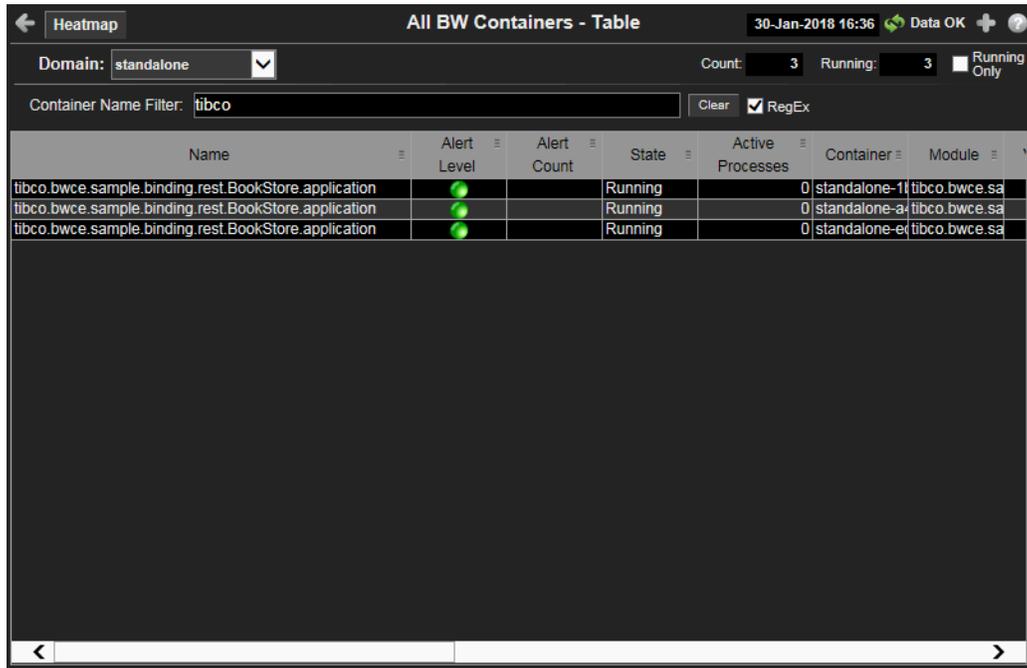
<b>Failed Count</b>	The total number of failed processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Created / sec</b>	The number of processes created per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Suspended / sec</b>	The number of suspended processes per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Failed / sec</b>	The number of failed processes per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Exec Time / sec</b>	The process execution time per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Most Recent Exec Time</b>	The execution time for the most recently executed process in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Average Exec Time</b>	The average execution time for all processes in the heatmap rectangle, calculated by dividing the delta execution time for the interval by the delta completed, or the number of process instances that completed in the interval. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Average Elapsed Time</b>	The average elapsed time for all processes in the heatmap rectangle, calculated by dividing the delta elapsed time for the interval by the delta completed, or the number of process instances that completed in the interval. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

## All Containers Table

This display provides a view of the most critical BusinessWorks container alert states pertaining to process creation and execution for all nodes on which the containers are deployed in a tabular format. Use this display to quickly identify containers with critical alerts. Each row in the table is a container in the selected domain.

Select a domain from the drop-down menu to view associated containers and, optionally, enter a string in the **Container Name Filter** field to further limit the list of containers shown in the display. You can click a column header to sort column data in numerical or alphabetical order.

To view additional details for a specific container, drill-down and investigate by clicking the row in the table for the desired container, which opens the “Single Container Summary” display.



**Title Bar** (possible features are):

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu, Table open commonly accessed displays.
- 6,047 The number of items currently in the display.
- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

**Filter By:**

The display might include these filtering options:

- Domain:** Choose a domain to show data for in the display.
- Container Name Filter** Enter a string (all or part of a container name) to filter the data shown in the display. If you enter part of a container name, you must enter "\*" before and/or after the string. For example, if you have a container named ContNameOne, you could filter using \*Name\*, \*NameOne, or ContName\*.
- Clear** Clears the **Container Name Filter** entries from the display.
- RegEx** Toggles the **Container Name Filter** to accept Regular Expressions for filtering. For example, if your application name is ContNameOne and this option was toggled on, you could enter "Name" (without using "\*" to display the container in the table).

**Fields and Data:**

- Count:** The total number of containers listed in the table.

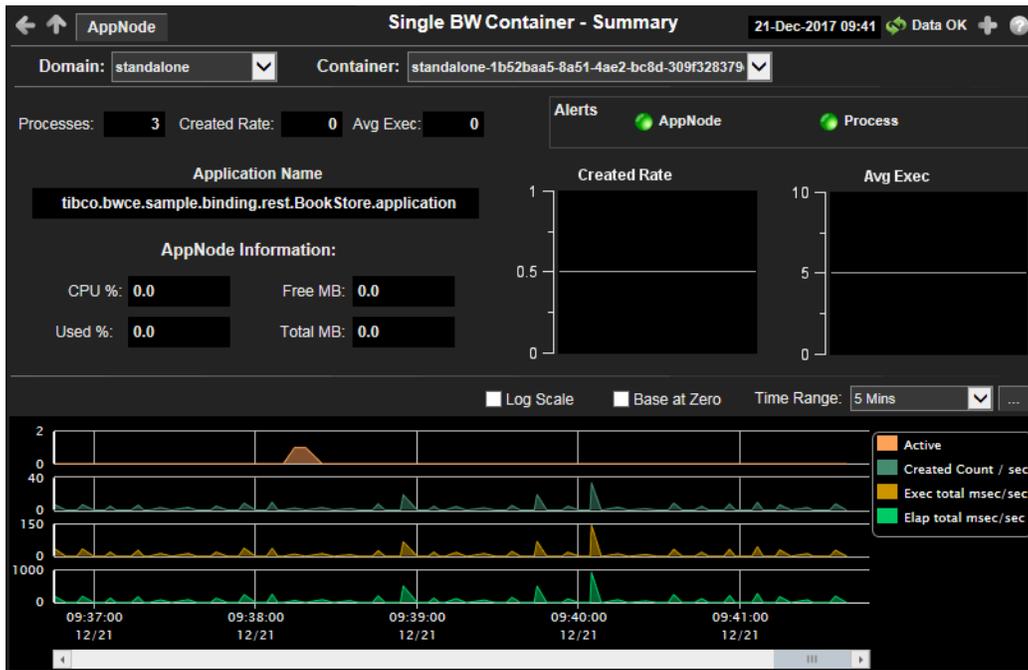
- Running** The total number of containers that are currently running.
- Running Only** Select to show only running containers in the display.

**Table:**

Each row in the table is a different application.

<b>Name</b>	The name of the container.
<b>Alert Level</b>	The most critical alert state for alerts in the row:  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
<b>Alert Count</b>	The total number of active alerts in the container.
<b>State</b>	The current status of the application. Valid values are <b>Running</b> and <b>Stopped</b> .
<b>Active Processes</b>	The number of currently active processes in the container.
<b>Container</b>	The name of the container.
<b>Module</b>	The name of the container module.
<b>Version</b>	The version of the container.
<b>Average Exec Time</b>	The average number of seconds for all processes to execute.
<b>Completed Processes</b>	The number of completed processes in the container.
<b>Expired</b>	When checked, data has not been received from this host in the specified amount of time.
<b>Failed Processes</b>	The number of failed processes in the container.
<b>Recent Exec Time</b>	The number of seconds for the most recently executed process.
<b>Created /sec</b>	The number of processes created per second in the container.
<b>Failed /sec</b>	The number of process failures per second in the container.
<b>Suspended / sec</b>	The number of process suspensions per second in the container.
<b>Exec Time/ sec</b>	The number of processes executed per second in the container.
<b>Suspended Processes</b>	The number of suspended application processes in the container.
<b>Time Stamp</b>	The date and time the row data was last updated.
<b>Source</b>	Name of RTView Data Server sending this data (or localhost).

## Single Container Summary



**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** , **Table** open commonly accessed displays.
- 6,047** The number of items currently in the display.

- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

**Filter By:**

The display might include these filtering options:

- Domain:** Select the domain for which you want to view data in the display.
- Container** Choose the container for which you want to view data in the display.

**Fields and Data:**

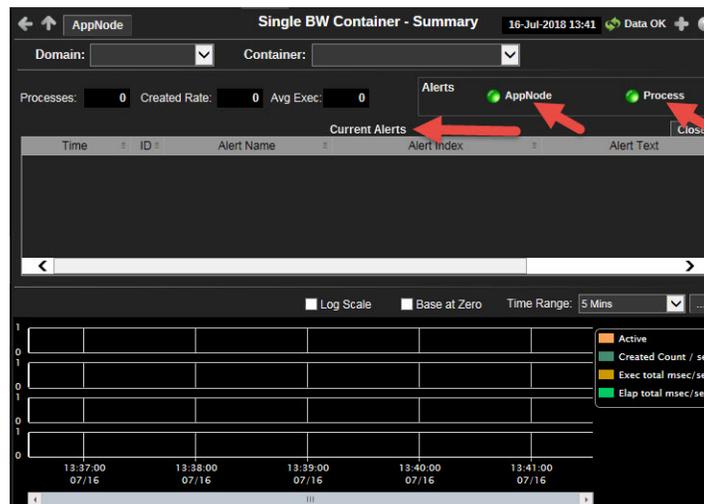
- Processes:** The number of processes currently running on the selected container.
- Created Rate:** The number of processes created per second on the selected container.
- Avg Exec:** The average number of seconds for processes to execute on the selected container.

### Alerts

Indicates the greatest severity level and the number of open **AppNode** and **Process** alerts for the selected container. Values range from **0** to **2**, where **2** is the greatest Severity:

- One or more alerts exceeded their ALARM LEVEL threshold.
- One or more alerts exceeded their WARNING LEVEL threshold.
- No alert thresholds have been exceeded.

Click on the alert indicator to display a table listing the current alerts for the selected container. Click the **Close** button (for the current alerts table) to close the table.



**Application Name** The name of the application running on the container.

### AppNode Information

**CPU %** The percentage of CPU used by the AppNode.

**Used %** The percentage of memory used by the AppNode.

**Free MB** The amount of free memory, in megabytes.

**Total MB** The total amount of used and free memory, in megabytes.

### Created Rate Bar Graph

The bar graph shows the current process creation rate per AppNode. Click to drill-down and investigate in the ["BW Single AppSlice Summary"](#) display.

### AvgExec Bar Graph

The bar graph shows the current average process execution rate per AppNode for the selected application. Click to drill-down and investigate in the ["BW Single AppSlice Summary"](#) display.

### Trend Graphs

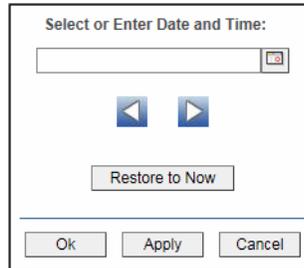
Traces the sum of process metrics across all processes in all slices of the selected container.

**Active** Traces the number of currently active application processes on the container.

**Created Count / sec** Traces the number of created application processes on the container.

**Exec total msec/sec** Traces the rate at which the application is accumulating process execution time, in milliseconds per second, on the container.

- Elap total msec/sec** Traces the rate at which the application accumulates process elapsed time, in milliseconds per second, on the container.
- Log Scale** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Base at Zero** Select to use zero (**0**) as the Y axis minimum for all graph traces.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## BW AppNodes

These displays present internal JVM memory and host CPU utilization for BusinessWorks AppNodes and their resources. This is useful because the AppNode performance is dependent on both internal and external factors and they sometimes interact. Displays in this View are:

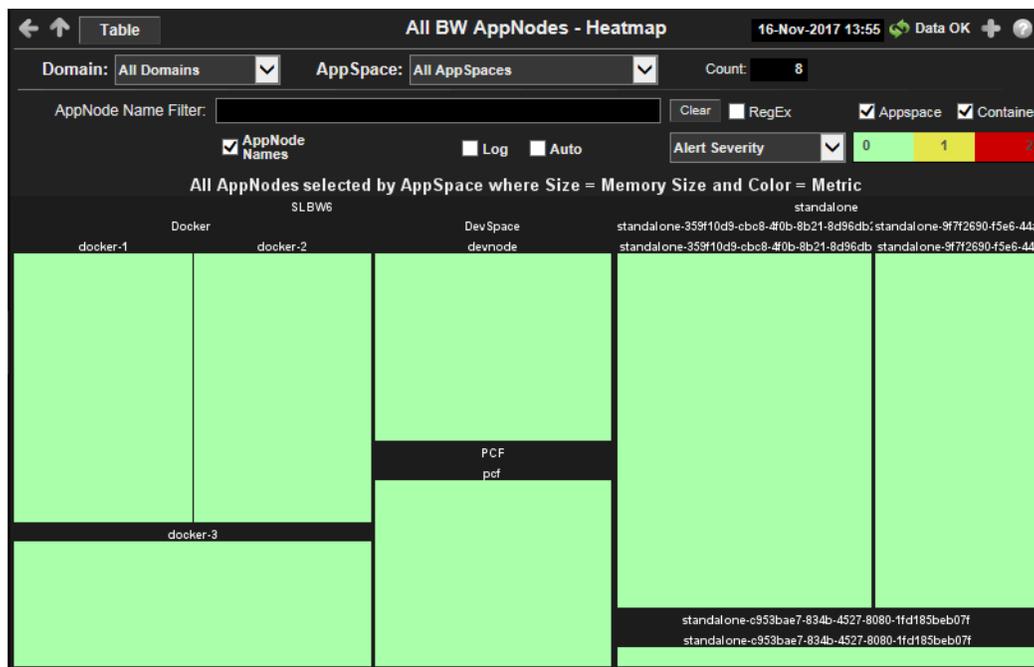
- ["BW All AppNodes Heatmap" on page 1265](#): A color-coded heatmap view of utilization metrics.
- ["BW All AppNodes Table" on page 1268](#): A tabular view of all available utilization data in this BusinessWorks View.
- ["BW Single AppNode Summary" on page 1270](#): Current and historical metrics for a single AppNode.

### BW All AppNodes Heatmap

View the most critical JVM memory and host resource utilization for BusinessWorks AppNodes. Use this display to quickly identify AppNodes with critical alerts.

Each rectangle in the heatmap represents an AppNode. The rectangle color indicates the most critical alert state associated with the AppNode. The rectangle size represents the maximum memory used in the rectangle; a larger size is a larger value.

Choose a domain and AppSpace from the drop-down menus. Choose a different metric to display from the **Metric** drop-down menu. Enter a string in the **AppNode Name Filter** field to limit data shown in the display. Use the **AppNode Names** check-box  to include or exclude labels in the heatmap. Mouse over a rectangle to see additional metrics. By default, this display shows **Alert Severity**. Drill-down and investigate by clicking a rectangle in the heatmap to view details for the selected application in the "BW Single AppNode Summary" display.



#### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** **Table** open commonly accessed displays.
- 6,047** The number of items currently in the display.

- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

#### Filter By:

The display might include these filtering options:

- Domain:** Choose a domain to show data for in the display.
- AppSpace** Choose an AppSpace to show data for in the display.
- Count:** The total number of AppNodes in the AppSpace.

<b>AppNode Name Filter</b>	Enter a string to limit data shown in the display.
	<b>Clear</b> Clears the <b>Application Name Filter</b> entries from the display.
<b>RegEx</b>	Toggles the <b>Search Text</b> field to accept Regular Expressions for filtering.
<b>AppSpace</b>	When selected, those AppNodes deployed in an AppSpace display in the heatmap.
<b>Container</b>	When selected, those AppNodes deployed in a container display in the heatmap.
<b>AppNode Names</b>	Check to include labels in the heatmap.
<b>Log</b>	Select to enable a logarithmic scale. Use <b>Log Scale</b> to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. <b>Log Scale</b> makes data on both scales visible by applying logarithmic values rather than actual values to the data.
<b>Auto</b>	Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value. NOTE: Some metrics auto-scale automatically, even when <b>Auto</b> is not selected.
<b>Metric</b>	Choose a metric to view in the display.
<b>Alert Severity</b>	<p>The maximum level of alerts in the heatmap rectangle. Values range from <b>0</b> - <b>2</b>, as indicated in the color gradient  bar, where <b>2</b> is the highest Alert Severity:</p> <ul style="list-style-type: none"> <li> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Alert Count</b>	The total number of critical and warning alerts in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.
<b>CPU Used%</b>	The percent (%) CPU used in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Memory Used%</b>	The percent (%) memory used in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Active Processes</b>	The number of currently active processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

**Created Processes** The number of processes created in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

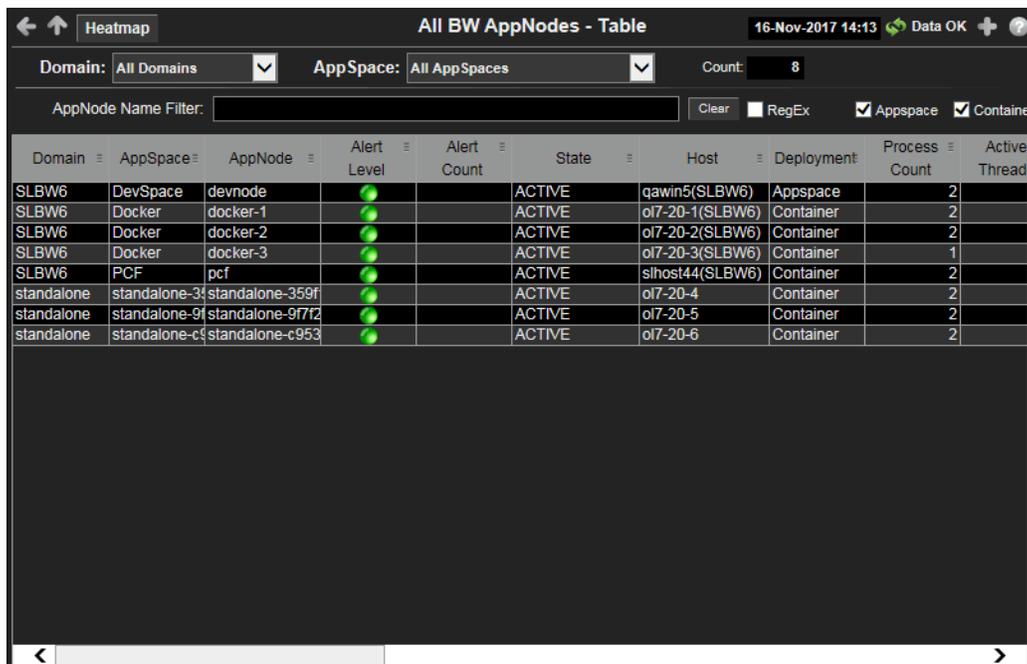
**Created /sec** The number of processes created per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

### BW All AppNodes Table

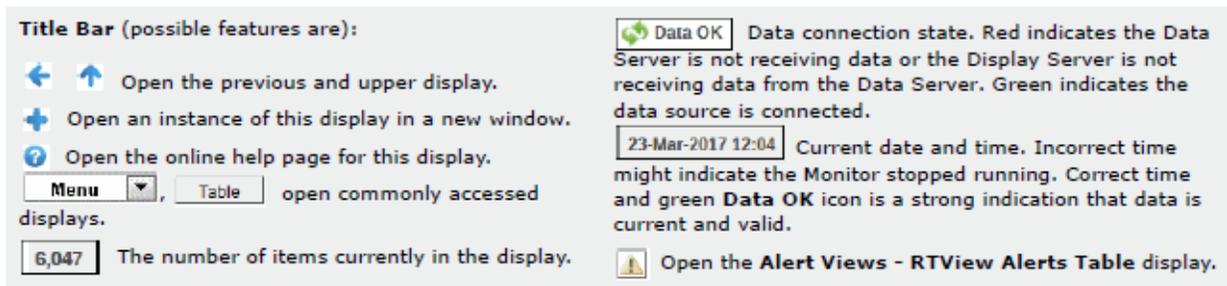
View BusinessWorks data shown in the "BW All AppNodes Heatmap" display, and additional details, in a tabular format. Use this display to view all available data for each AppNode by Domain and AppSpace.

Each row in the table is an AppNode. Choose a domain and AppSpace from the drop-down menus. Click a column header to sort column data in numerical or alphabetical order. Enter a string in the **Application Name Filter** field to limit data shown in the display.

Drill-down and investigate by clicking a row to view details for the selected AppNode in the "BW Single AppNode Summary" display.



Domain	AppSpace	AppNode	Alert Level	Alert Count	State	Host	Deployment	Process Count	Active Thread
SLBW6	DevSpace	devnode			ACTIVE	qawin5(SLBW6)	Appspace	2	
SLBW6	Docker	docker-1			ACTIVE	oi7-20-1(SLBW6)	Container	2	
SLBW6	Docker	docker-2			ACTIVE	oi7-20-2(SLBW6)	Container	2	
SLBW6	Docker	docker-3			ACTIVE	oi7-20-3(SLBW6)	Container	1	
SLBW6	PCF	pcf			ACTIVE	sihost44(SLBW6)	Container	2	
standalone	standalone-3	standalone-359f			ACTIVE	oi7-20-4	Container	2	
standalone	standalone-9	standalone-9772			ACTIVE	oi7-20-5	Container	2	
standalone	standalone-c	standalone-c953			ACTIVE	oi7-20-6	Container	2	

**Filter By:**

The display might include these filtering options:

<b>Domain:</b>	Choose a domain to show data for in the display.
<b>AppSpace</b>	Choose an AppSpace to show data for in the display.
<b>Count:</b>	The total number of rows in the table.
<b>AppNode Name Filter</b>	Enter a string to limit data shown in the display.
<b>Clear</b>	Clears the <b>Application Name Filter</b> entries from the display.
<b>RegEx</b>	Toggles the <b>Search Text</b> field to accept Regular Expressions for filtering.
<b>AppSpace</b>	When selected, those AppNodes deployed in an AppSpace display in the AppNodes table.
<b>Container</b>	When selected, those AppNodes deployed in a container display in the AppNodes table.

**Count:** The total number of rows in the table.

**Table:**

Column values describe the AppNode.

<b>Domain</b>	The domain in which the AppNode resides.
<b>AppSpace</b>	The AppSpace in which the AppNode resides.
<b>AppNode</b>	The name of the AppNode.
<b>Alert Level</b>	The most critical alert state for alerts in the row: <ul style="list-style-type: none"> <li> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Alert Count</b>	The total number of active alerts for the AppNode.
<b>Host</b>	The host on which the AppNode resides.
<b>Process Count</b>	The number of processes running.
<b>Active Threads</b>	The number of currently active threads.
<b>Total Memory</b>	The total amount of used and free memory, in megabytes.
<b>Used Memory</b>	The amount of used memory, in megabytes.

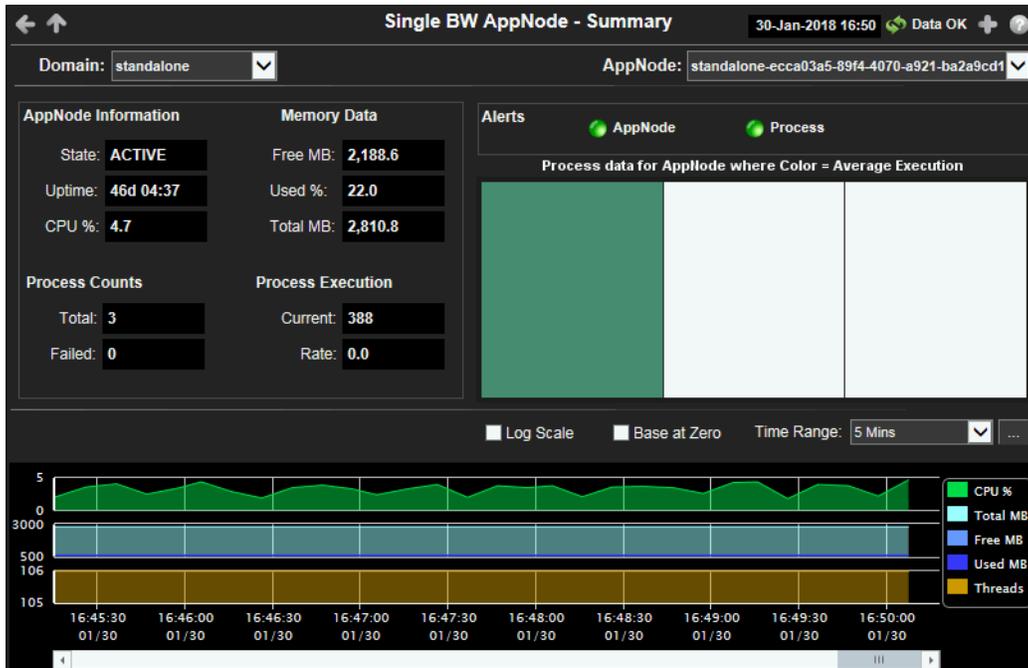
<b>Free Memory</b>	The amount of free memory, in megabytes.
<b>Used Memory%</b>	The percent (%) used memory.
<b>Used CPU%</b>	The percent (%) used CPU.
<b>System Process ID</b>	A unique string identifier for the process.
<b>Up Since</b>	The date and time the AppNode was last started.
<b>Active Processes</b>	The number of currently active processes.
<b>Suspended Processes</b>	The number of suspended application processes.
<b>Failed Processes</b>	The number of failed application processes.
<b>Completed Processes</b>	The number of completed application processes.
<b>Created /sec</b>	The number of application processes created per second.
<b>Suspended / sec</b>	The number of application processes suspended per second.
<b>Failed /sec</b>	The number of failed application processes per second.
<b>Exec Time / sec</b>	The number of application processes executed per second.
<b>Recent Exec Time</b>	The number of seconds for the most recently executed process.
<b>Average Exec Time</b>	The average number of seconds for all processes to execute.
<b>Time Stamp</b>	The date and time the row data was last updated.
<b>Source</b>	Name of RTView Data Server sending this data (or localhost).

## BW Single AppNode Summary

View current and historical utilization and performance metrics for a single BusinessWorks AppNode. Use this display to investigate performance issues on an AppNode.

This display includes a heatmap showing most critical alerts pertaining to process execution, and trend graphs tracing CPU utilization and thread count.

Choose a domain, AppSpace and AppNode from the drop-down menus. Use the **Time-Range** to "zoom-in" or "zoom-out" on a specific time frame in the trend graph. Drill-down and investigate by clicking an AppNode in the table to view details in the ["BW Single AppNode Summary"](#) display.



**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** **Table** open commonly accessed displays.
- 6,047** The number of items currently in the display.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

Open the **Alert Views - RTView Alerts Table** display.

**Filter By:**

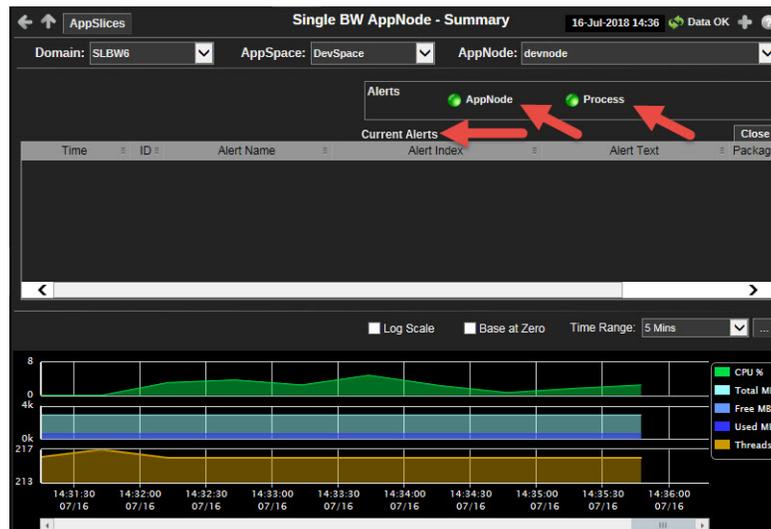
The display might include these filtering options:

- Domain:** Choose a domain to show data for in the display.
- AppSpace** Choose an AppSpace to show data for in the display.
- AppNode:** Choose an AppNode to show data for in the display.

**Fields and Data:**

- AppNode Information**
  - Uptime:** The number of days, hours and minutes since the AppNode started.
  - CPU%** The percent (%) CPU used on the AppNode.
  - Threads:** The number of currently active threads for the AppNode.
- Memory Data**
  - Free:** The amount of available memory on the AppNode.
  - Used%** The percent (%) memory used on the AppNode.

	<b>Total</b>	The total amount of memory on the AppNode.
<b>Process Counts</b>	<b>Total:</b>	The number of currently active processes for the AppNode.
	<b>Failed:</b>	The number of failed processes for the AppNode.
<b>Process Execution</b>	<b>Current</b>	The number of processes executed by the AppNode.
	<b>Rate:</b>	The number of processes executed per second.
<b>Alerts</b>	Indicates the greatest severity level and the number of open <b>AppNode</b> and <b>Process</b> alerts for the selected AppNode. Values range from <b>0</b> to <b>2</b> , where <b>2</b> is the greatest Severity:	
		One or more alerts exceeded their ALARM LEVEL threshold.
		One or more alerts exceeded their WARNING LEVEL threshold.
		No alert thresholds have been exceeded.
	Click on the alert indicator to display a table listing the current alerts for the selected AppNode. Click the <b>Close</b> button (for the current alerts table) to close the table.	



### Heatmap

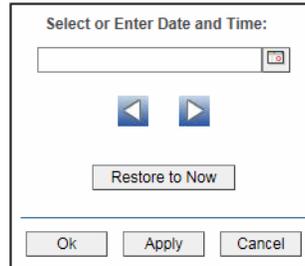
Each rectangle in the heatmap represents an AppSlice. The rectangle color indicates the most critical **Average Exec Time** alert state associated with the AppSlice. The rectangle size represents the maximum number of processes executed in the rectangle; a larger size is a larger value. Click a rectangle to drill-down and investigate in the ["BW Single AppSlice Summary"](#) display.

### Trend Graphs

Traces the sum of process metrics across all processes for all applications on the AppNode.

- **CPU%:** The percent (%) CPU used on the AppNode.
- **Total MB:** The amount of memory used.
- **Free MB:** The amount of available memory.
- **Used MB:** The amount of used memory.
- **Threads:** The number of threads.

- Log Scale** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Base at Zero** Select to use zero (**0**) as the Y axis minimum for all graph traces.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## BW AppSlices

These displays present process metrics totaled by Application and AppNode for AppSlices. This is useful to see how the application is distributed and how each part of it is performing. The AppSlice is the part of an application running on a specific AppNode when the application is deployed to multiple AppNodes. Displays in this View are:

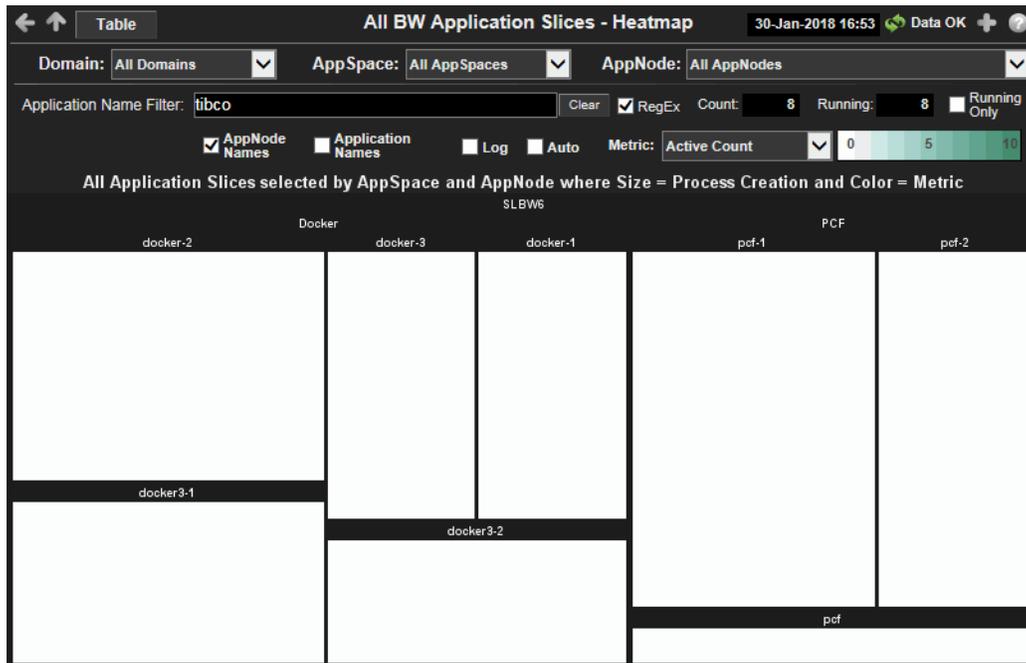
- ["BW All AppSlices Heatmap" on page 1273](#): A color-coded heatmap view of process creation and execution metrics.
- ["BW All AppSlices Table" on page 1276](#): A tabular view of all available data in this BusinessWorks View.
- ["BW Single AppSlice Summary" on page 1278](#): Current and historical metrics for a single AppSlice.

### BW All AppSlices Heatmap

View the most critical performance metrics for BusinessWorks AppSlices. Use this display to quickly identify AppSlices with high process execution numbers.

Each rectangle in the heatmap represents an AppSlice. The rectangle color indicates the process execution numbers for the AppSlice. The rectangle size represents the number of processes created in the rectangle; a larger size is a larger value.

Choose a domain, AppSpace and AppNode from the drop-down menus. Choose a different metric to display from the **Metric** drop-down menu. Enter a string in the **Application Name Filter** field to limit data shown in the display. Use the **AppNode Names** and **Application Names** check-boxes  to include or exclude labels in the heatmap. Mouse over a rectangle to see additional metrics. By default, this display shows **Active Count**. Drill-down and investigate by clicking a rectangle in the heatmap to view details for the selected application in the "BW Single AppSlice Summary" display.



#### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** , **Table** open commonly accessed displays.
- 6,047** The number of items currently in the display.

- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

#### Filter By:

The display might include these filtering options:

- Domain:** Choose a domain to show data for in the display.
- AppSpace** Choose an AppSpace to show data for in the display.
- AppNode:** Choose an AppNode to show data for in the display.

#### Fields and Data:

- Application Name Filter** Enter a string to limit data shown in the display.

	<b>Clear</b>	Clears the <b>Application Name Filter</b> entries from the display.
<b>RegEx</b>		Toggles the <b>Search Text</b> field to accept Regular Expressions for filtering.
<b>Count</b>		The number of AppNodes in the display.
<b>Running</b>		The total number of AppSpaces currently running in the display.
<b>Running Only</b>		Select to show only running applications in the display.
<b>AppNode Names</b>		Check to include labels in the heatmap.
<b>Application Names</b>		Check to include labels in the heatmap.
<b>Log</b>		Select to enable a logarithmic scale. Use <b>Log Scale</b> to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. <b>Log Scale</b> makes data on both scales visible by applying logarithmic values rather than actual values to the data.
<b>Auto</b>		Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value. NOTE: Some metrics auto-scale automatically, even when <b>Auto</b> is not selected.
<b>Metric</b>		Choose a metric to view in the display.
	<b>Active Count</b>	The total number of active processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
	<b>Completed Count</b>	The total number of completed processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
	<b>Suspended Count</b>	The total number of suspended processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
	<b>Failed Count</b>	The total number of failed processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
	<b>Created / sec</b>	The number of processes created per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
	<b>Suspended / sec</b>	The number of suspended processes per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

<b>Failed / sec</b>	The number of failed processes per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Exec Time / sec</b>	The process execution time per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Most Recent Exec Time</b>	The execution time for the most recently executed process in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Average Exec Time</b>	The average execution time for all processes in the heatmap rectangle, calculated by dividing the delta execution time for the interval by the delta completed, or the number of process instances that completed in the interval. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Average Elapsed Time</b>	The average elapsed time for all processes in the heatmap rectangle, calculated by dividing the delta elapsed time for the interval by the delta completed, or the number of process instances that completed in the interval. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

## BW All AppSlices Table

View BusinessWorks data shown in the ["BW All AppSlices Heatmap"](#), and additional details, in a tabular format.

Each row in the table is an AppSlice. Choose a domain (or **All Domains**), an AppSpace (or **All AppSpaces**) and an AppNode (or **All AppNodes**) from the drop-down menus. Click a column header to sort column data in numerical or alphabetical order. Enter a string in the **Application Name Filter** field to limit data shown in the display.

Drill-down and investigate by clicking a row to view details in the ["BW Single AppSlice Summary"](#) display.

Domain	AppSpace	AppNode	Name	Version	State
SLBW6	Docker	docker-1	tibco.bwce.sample.binding.rest.BookStore.application	1.0	Running
SLBW6	Docker	docker-2	tibco.bwce.sample.binding.rest.BookStore.application	1.0	Running
SLBW6	Docker	docker-3	tibco.bwce.sample.binding.rest.BookStore.application	1.0	Running
SLBW6	Docker	docker3-1	tibco.bwce.sample.binding.rest.BookStore.application	1.0	Running
SLBW6	Docker	docker3-2	tibco.bwce.sample.binding.rest.BookStore.application	1.0	Running
SLBW6	PCF	pcf	tibco.bwce.sample.binding.rest.BookStore.application	1.0	Running
SLBW6	PCF	pcf-1	tibco.bwce.sample.binding.rest.BookStore.application	1.0	Running
SLBW6	PCF	pcf-2	tibco.bwce.sample.binding.rest.BookStore.application	1.0	Running

**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** **Table** open commonly accessed displays.
- 6,047** The number of items currently in the display.
- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

**Filter By:**

The display might include these filtering options:

- Domain:** Choose a domain to show data for in the display.
- AppSpace** Choose an AppSpace to show data for in the display.
- AppNode** Choose an AppNode to show data for in the display.
- Application Name Filter** Enter a string to limit data shown in the display.
- Clear** Clears the **Application Name Filter** entries from the display.
- RegEx** Toggles the **Application Name Filter** to accept Regular Expressions for filtering.

**Fields and Data:**

- Count:** The total number of rows in the table.
- Running** The total number of applications currently running in the AppSpace.

**Running Only** Select to show only running applications in the display.

**Table:**

Each row in the table is a different AppNode.

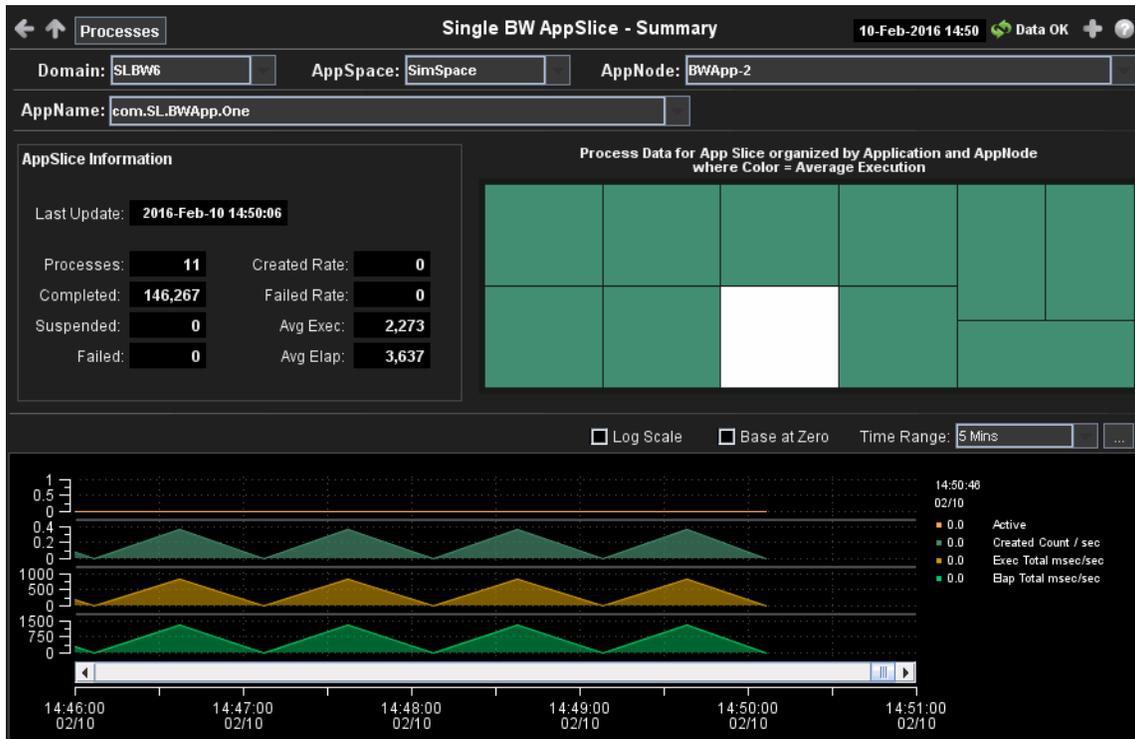
<b>Domain</b>	The domain in which the AppSpace resides.
<b>AppSpace</b>	The AppSpace the AppNode is associated with.
<b>AppNode</b>	The name of the selected AppNode.
<b>Name</b>	The name of the application.
<b>Version</b>	The application version.
<b>State</b>	The current status of the application. Valid values are <b>Running</b> and <b>Stopped</b> .
<b>Module</b>	The application module.
<b>Shared Module</b>	The shared module, if any.
<b>Active Processes</b>	The number of currently active application processes.
<b>Suspended Processes</b>	The number of suspended application processes.
<b>Failed Processes</b>	The number of failed application processes.
<b>Completed Processes</b>	The number of completed application processes.
<b>Created /sec</b>	The number of application processes created per second.
<b>Suspended / sec</b>	The number of application process suspensions per second.
<b>Failed /sec</b>	The number of application process failures per second.
<b>Exec Time / sec</b>	The number of processes executed per second.
<b>Recent Exec Time /sec</b>	The number of seconds for the most recently executed process.
<b>Average Exec Time</b>	The average number of seconds for all processes to execute.
<b>Time Stamp</b>	The date and time the row data was last updated.
<b>Source</b>	Name of RTView Data Server sending this data (or localhost).

## BW Single AppSlice Summary

View current and historical utilization and performance metrics for a single BusinessWorks AppSlice. Use this display to investigate performance issues on an AppSlice level.

This display includes a heatmap showing the most critical process execution alerts for AppSlices on the selected application, and trend graphs tracing process execution times.

Choose a domain, AppSpace, AppNode and AppNode from the drop-down menus. Use the **Time-Range** to "zoom-in" or "zoom-out" on a specific time frame in the trend graph. Drill-down and investigate by clicking a process in the heatmap to view details in the "BW Single Process Summary" display.



**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- 6,047** The number of items currently in the display.
- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

**Filter By:**

The display might include these filtering options:

- Domain:** Choose a domain to show data for in the display.
- AppSpace** Choose an AppSpace to show data for in the display.
- AppNode:** Choose an AppNode to show data for in the display.
- AppName:** Choose an AppName to show data for in the display.

**Fields and Data:**

- AppSlice Information** **Last Update:** The date and time the data was last updated.

<b>Processes</b>	The number of active processes.
<b>Completed:</b>	The total number of completed processes summed across all processes in one AppSlice of the application.
<b>Suspended:</b>	The total number of suspended processes
<b>Failed:</b>	The total number of failed processes
<b>Created Rate:</b>	The number of application processes created per second.
<b>Failed Rate:</b>	The number of failed application processes per second.
<b>Avg Exec:</b>	The average number of seconds for processes to execute.
<b>Avg Elap:</b>	The average amount of elapsed time for processes, in seconds.

### Heatmap

Each rectangle in the heatmap represents one process in an AppSlice. The rectangle color indicates the most critical **Average Exec Time** alert state associated with the AppSlice. The rectangle size represents the processes execution time in the rectangle; a larger size is a larger value. Click a rectangle to drill-down and investigate in the "[BW Single Process Summary](#)" display.

### Trend Graphs

Traces the sum across all processes in one AppSlice of the application.

- **Active:** Traces the number of active processes.
- **Created Count:** Traces the number of processes created.
- **Exec Total msec/sec:** Traces the rate at which the application accumulates process execution time, in milliseconds per second.
- **Elap Total msec/sec:** Traces the rate at which the application is accumulating process elapsed time, in milliseconds per second.

**Log Scale** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.

**Base at Zero** Select to use zero (0) as the Y axis minimum for all graph traces.

**Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .

By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## BW Processes

These displays present performance data for BusinessWorks processes. Use these displays to verify that individual BusinessWorks processes are executing and using resources as expected. Displays in this View are:

- [“BW All Processes Heatmap” on page 1281](#): A color-coded heatmap view of selected process performance metrics.
- [“BW All Processes Table” on page 1284](#): A tabular view of all available process performance data in this BusinessWorks View.
- [“BW Single Process Summary” on page 1287](#): Current and historical metrics for a single process.

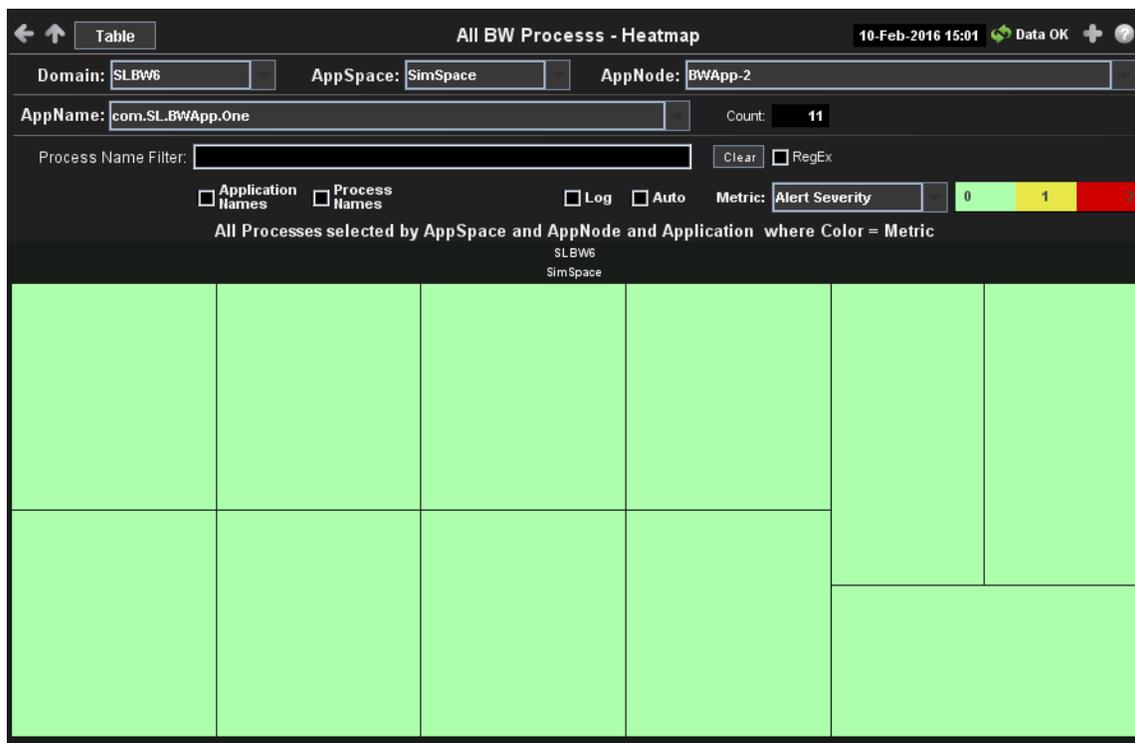
### BW All Processes Heatmap

View the most critical BusinessWorks alerts pertaining to process creation and execution. Use this display to quickly identify processes with critical alerts.

Each rectangle in the heatmap represents a process. The rectangle color indicates the most critical alert state associated with the processes (the rectangle size is uniform for all processes.)

Choose a domain, applications, AppNode and AppSpace from the drop-down menus. Choose a different metric to display from the **Metric** drop-down menu. Enter a string in the **Process Name Filter** field to limit data shown in the display. Use the **Application Names** and **Process Names** check-boxes  to include or exclude labels in the heatmap. Mouse over a rectangle to see additional metrics. By default, this display shows **Alert Severity**.

Drill-down and investigate by clicking a rectangle in the heatmap to view details in the [“BW Single Process Summary”](#) display.



**Title Bar** (possible features are):

-   Open the previous and upper display.
-  Open an instance of this display in a new window.
-  Open the online help page for this display.
- open commonly accessed displays.
- The number of items currently in the display.

 **Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

 Open the **Alert Views - RTView Alerts Table** display.

**Filter By:**

The display might include these filtering options:

- Domain:** Choose a domain to show data for in the display.
- AppSpace** Choose an AppSpace to show data for in the display.
- AppNode:** Choose an AppNode to show data for in the display.
- AppName** Choose an AppName to show data for in the display.
- Count:** The total number of processes currently shown in the display.

**Fields and Data:**

- Process Name Filter** Enter a string to limit data shown in the display.
  - Clear** Clears the **Processes Name Filter** entries from the display.
- RegEx** Toggles the **Processes Name Filter** to accept Regular Expressions for filtering.
- Application Names** Check to include labels in the heatmap.
- Process Names** Check to include labels in the heatmap.
- Log** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Auto** Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value. NOTE: Some metrics auto-scale automatically, even when **Auto** is not selected.
- Metric** Choose a metric to view in the display.
- Alert Severity** The maximum level of alerts in the heatmap rectangle. Values range from **0** - **2**, as indicated in the color gradient  bar, where **2** is the highest Alert Severity:
  -  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
  -  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
  -  Green indicates that no metrics have exceeded their alert thresholds.

<b>Alert Count</b>	The total number of critical and warning alerts in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.
<b>Active Count</b>	The total number of active processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Completed Count</b>	The total number of completed processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Suspended Count</b>	The total number of suspended processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Failed Count</b>	The total number of failed processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Created / sec</b>	The number of processes created per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Suspended / sec</b>	The number of suspended processes per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Failed / sec</b>	The number of failed processes per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Exec Time / sec</b>	The process execution time per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Most Recent Exec Time</b>	The execution time for the most recently executed process in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

**Average Exec Time**

The average execution time for all processes in the heatmap rectangle, calculated by dividing the delta execution time for the interval by the delta completed, or the number of process instances that completed in the interval. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

**Most Recent Elapsed Time**

The elapsed time for the most recent process in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

**Average Elapsed Time**

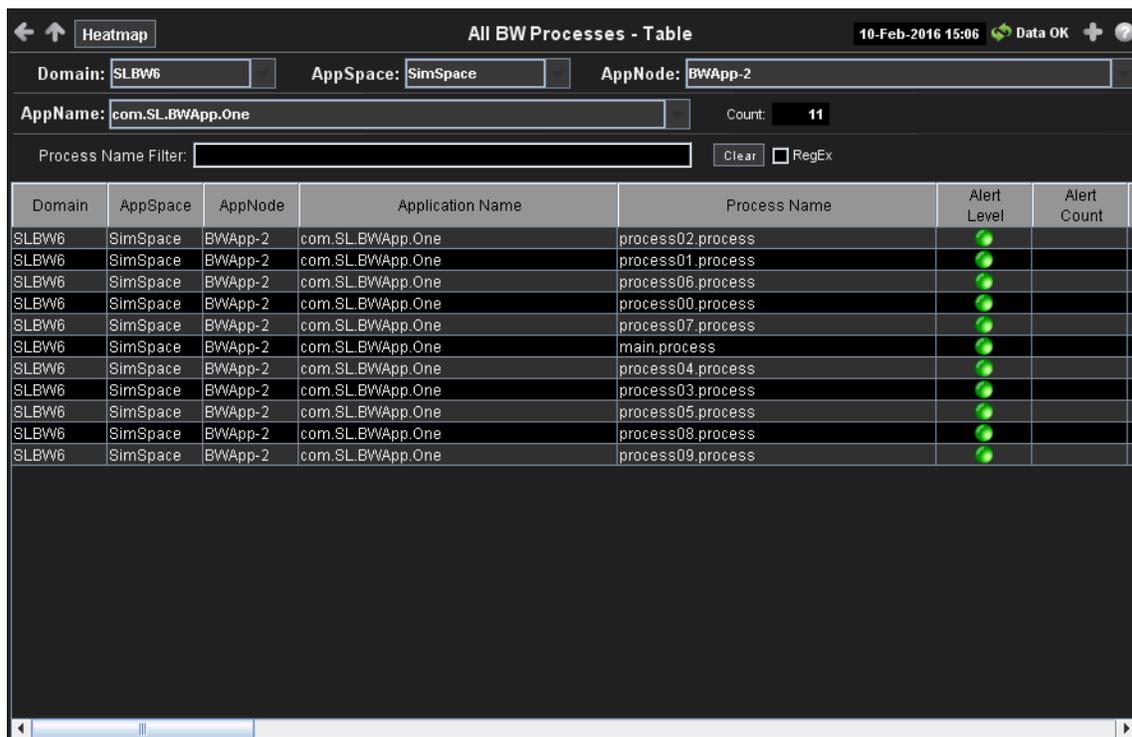
The average elapsed time for all processes in the heatmap rectangle, calculated by dividing the delta elapsed time for the interval by the delta completed, or the number of process instances that completed in the interval. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

**BW All Processes Table**

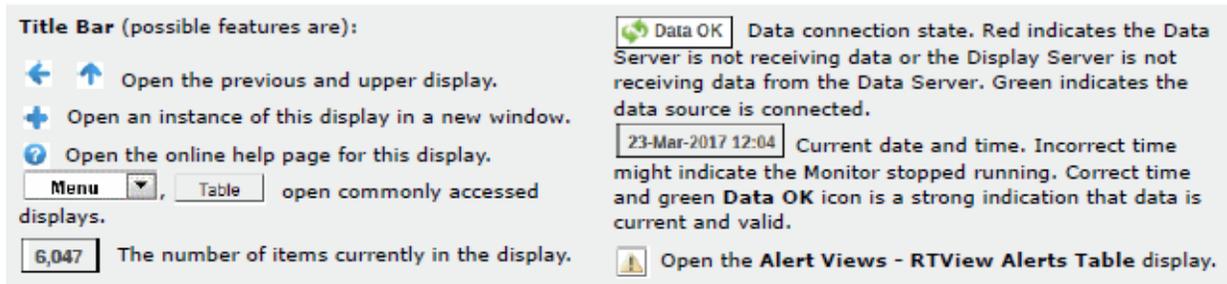
View BusinessWorks data shown in the "BW All Applications Heatmap", and additional details, in a tabular format.

Each row in the table is a process. Choose a domain, applications, AppNode and AppSpace from the drop-down menus. Click a column header to sort column data in numerical or alphabetical order. Enter a string in the **Process Name Filter** field to limit data shown in the display.

Drill-down and investigate by clicking a row to view details for the selected process in the "BW Single Process Summary" display



Domain	AppSpace	AppNode	Application Name	Process Name	Alert Level	Alert Count
SLBW6	SimSpace	BWApp-2	com.SL.BWApp.One	process02.process		
SLBW6	SimSpace	BWApp-2	com.SL.BWApp.One	process01.process		
SLBW6	SimSpace	BWApp-2	com.SL.BWApp.One	process06.process		
SLBW6	SimSpace	BWApp-2	com.SL.BWApp.One	process00.process		
SLBW6	SimSpace	BWApp-2	com.SL.BWApp.One	process07.process		
SLBW6	SimSpace	BWApp-2	com.SL.BWApp.One	main.process		
SLBW6	SimSpace	BWApp-2	com.SL.BWApp.One	process04.process		
SLBW6	SimSpace	BWApp-2	com.SL.BWApp.One	process03.process		
SLBW6	SimSpace	BWApp-2	com.SL.BWApp.One	process05.process		
SLBW6	SimSpace	BWApp-2	com.SL.BWApp.One	process08.process		
SLBW6	SimSpace	BWApp-2	com.SL.BWApp.One	process09.process		

**Filter By:**

The display might include these filtering options:

- Domain:** Choose a domain to show data for in the display.
- AppSpace** Choose an AppSpace to show data for in the display.
- AppNode:** Choose an AppNode to show data for in the display.
- AppName** Choose an AppName to show data for in the display.

**Fields and Data:**

- Count:** The total number of processes in the AppSpace.
- Process Name Filter** Enter a string to limit data shown in the display.
- Clear** Clears the **Application Name Filter** entries from the display.
- RegEx** Toggles the **Application Name Filter** to accept Regular Expressions for filtering.

**Table:**

Each row in the table is a different AppSlice. Column values are associated with the process.

- Domain** The domain in which the process resides.
- AppSpace** The AppSpace in which the process resides.
- AppNode** The AppSpace in which the process resides.
- Application Name** The name of the application in which the process is running.
- Process Name** The name of the process.
- Alert Level** The most critical alert state for alerts in the row:
  - Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
  - Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
  - Green indicates that no metrics have exceeded their alert thresholds.
- Alert Count** The total number of active alerts for the process.
- Total Exec Time** Total execution time (in milliseconds) for all successfully completed process instances.
- Delta Exec Time** Execution time accumulated during the current polling period.

<b>Exec Time/ sec</b>	Delta execution time per second.
<b>Recent Exec Time</b>	Execution time (in milliseconds) of the most recently completed process instance.
<b>Total Elapsed Time</b>	Total elapsed time (in milliseconds) for all successfully completed process instances.
<b>Delta Elapsed Time</b>	Elapsed time accumulated during the current polling period.
<b>Elapsed Time/sec</b>	Delta elapsed time per second.
<b>Recent Elapsed Time</b>	Elapsed clock time (in milliseconds) of the most recently completed process instance.
<b>Active</b>	The number of currently active processes
<b>Created</b>	The number of processes created.
<b>Suspended</b>	The number of process suspensions.
<b>Failed</b>	The number of process failures.
<b>Completed</b>	The number of completed processes.
<b>Delta Active</b>	The number of active processes since the last data update.
<b>Active/sec</b>	The number of active processes per second.
<b>Delta Created</b>	The number of created processes since the last data update.
<b>Created/sec</b>	The number of created processes per second.
<b>Delta Suspended</b>	The number of suspended processes since the last data update.
<b>Suspended/ sec</b>	The number of suspended processes per second.
<b>Delta Completed</b>	The number of completed processes since the last data update.
<b>Completed/ sec</b>	The number of completed processes per second.
<b>Delta Failed</b>	The number of failed processes since the last data update.
<b>Failed/sec</b>	The number of failed processes per second.
<b>Min Exec Time</b>	Execution time (in milliseconds) of the process instance that has completed in the shortest amount of execution time.
<b>Max Exec Time</b>	Execution time (in milliseconds) of the process instance that has completed in the longest amount of execution time.
<b>Average Exec Time</b>	Average execution time (in milliseconds) for all successfully completed process instances.
<b>Min Elapsed Time</b>	Elapsed clock time (in milliseconds) of the process instance that has completed in the shortest amount of elapsed time.
<b>Max Elapsed Time</b>	Elapsed clock time (in milliseconds) of the process instance that has completed in the longest amount of elapsed time.

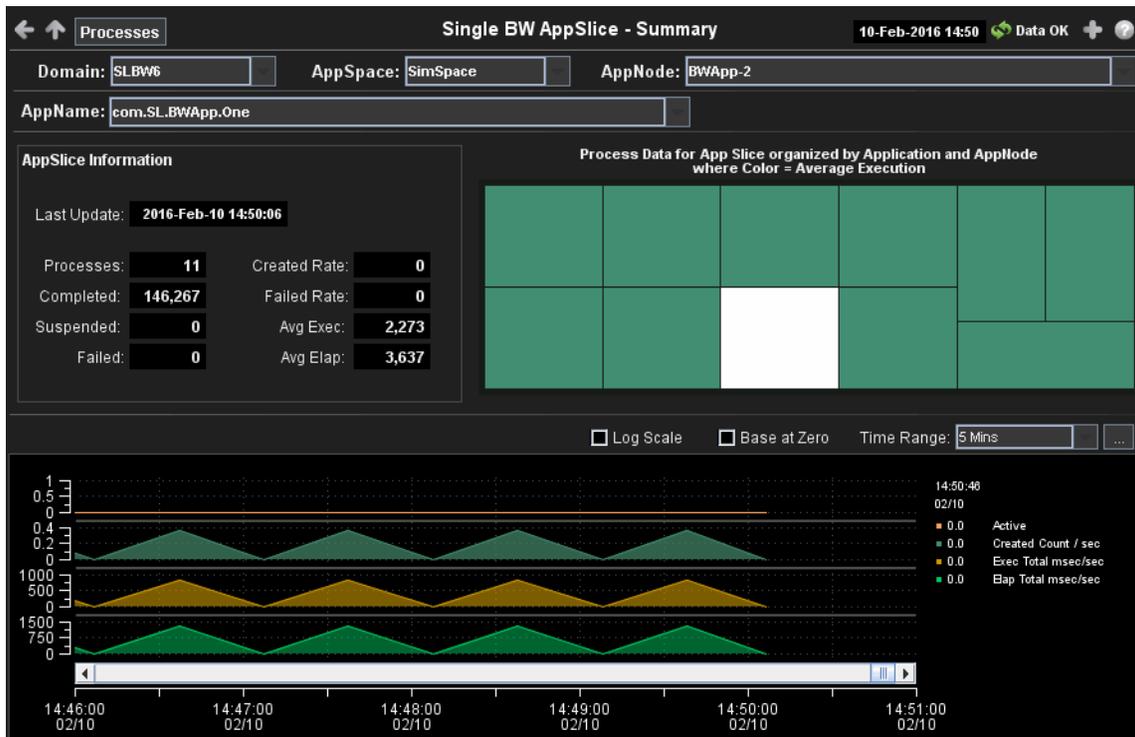
<b>Average Elapsed Time</b>	Average elapsed clock time (in milliseconds) for all successfully completed process instances.
<b>Count Since Reset</b>	The number of times the process has executed since statistics were reset.
<b>Main Process</b>	The name of the main process.
<b>Application Version</b>	The application version.
<b>Module Name</b>	The application module.
<b>Module Version</b>	The module version.
<b>Time Stamp</b>	The date and time the row data was last updated.
<b>Source</b>	Name of RTView Data Server sending this data (or localhost).

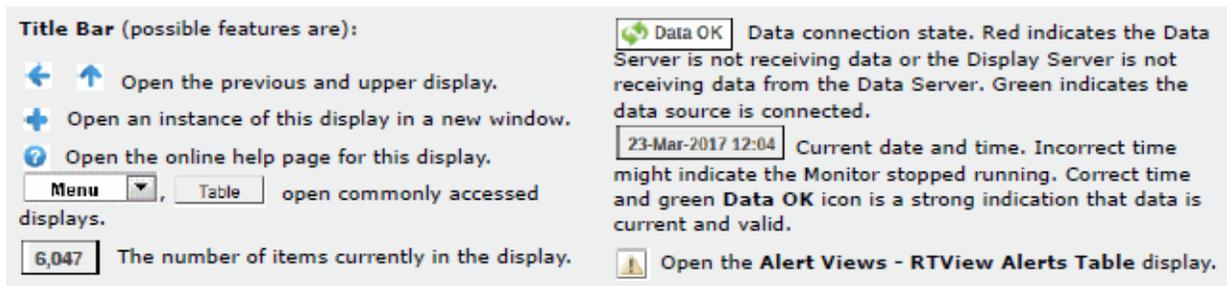
### BW Single Process Summary

View current and historical execution metrics for a single BusinessWorks process. Use this display to investigate performance issues for a process.

This display includes trend graphs tracing process and activity execution counts and times.

Choose a domain, application, AppNode and AppSpace from the drop-down menus. Use the **Time-Range** to "zoom-in" or "zoom-out" on a specific time frame in the trend graph.



**Filter By:**

The display might include these filtering options:

- Domain:** Choose a domain to show data for in the display.
- AppSpace** Choose an AppSpace to show data for in the display.
- AppNode:** Choose an AppNode to show data for in the display.
- AppName** Choose an application to show data for in the display.
- Process** Choose a process to show data for in the display.

**Fields and Data:**

- Activity Count:** The number of activities defined for the process.
- Main Process:** The name of the main process.
- Active** Number of active instances for this process definition. This number is calculated using the Hawk method named GetProcesses. This method returns information about process instances that are active at the time of update. The value here displays the current total count of all active instances discovered for this process definition. The trend below displays the same value over time.
- Active/sec** The number of currently active application processes per second.
- Created**
  - Total** The number of process instances created for this process definition.
  - Current** The number of process instances created this update cycle.
  - Rate** The number of process instances created per second.
- Completed**
  - Total** The number of process instances completed for this process definition.
  - Current** The number of process instances completed this update cycle.
  - Rate** The number of process instances completed per second.
- Errors**
  - Total** The number of errors accumulated by all process instances.
  - Current** The number of errors accumulated this update cycle.
  - Rate** The number of errors accumulated per second.
- Execution**
  - Min** The shortest execution time of any process instance, in milliseconds.
  - Max** The longest execution time of any process instance, in milliseconds.

	<b>Average</b>	The average execution time for all completed process instances, in milliseconds.
	<b>Current</b>	The amount of time accumulated this update cycle.
	<b>Rate</b>	The amount of time accumulated per second.
<b>Elapsed</b>	<b>Min</b>	The shortest elapsed time of any process instance, in milliseconds.
	<b>Max</b>	The longest elapsed time of any process instance, in milliseconds.
	<b>Average</b>	The average elapsed time for all completed process instances, in milliseconds.
	<b>Current</b>	The amount of elapsed time accumulated this update cycle.
	<b>Rate</b>	The amount of elapsed time accumulated per second.

### Trend Graphs

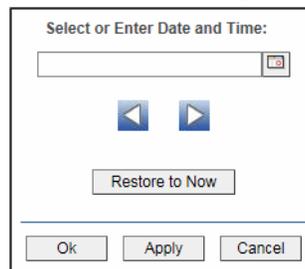
Traces application process and activity metrics for the selected process.

- **Active Count:** Traces the number of currently active processes.
- **Created Count:** Traces the number of created processes.
- **Process Elapsed Time/sec:** Traces the rate at which the application is accumulating process elapsed time, in milliseconds per second.
- **Process Exec Time/sec:** Traces the rate at which the application is accumulating process execution time, in milliseconds per second.
- **All Activities Exec Count/sec:** Traces the number of executed activities per second.
- **All Activities Exec Time/sec:** Traces the amount of execution time for executed activities per second.

**Log Scale** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.

**Base at Zero** Select to use zero (0) as the Y axis minimum for all graph traces.

**Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## BW5 Engines

These displays present performance metrics for BW5 Engines. Displays in this View are:

- ["All Engines Heatmap" on page 1290](#): Performance metrics of CPU and memory utilization for all BW Engines.
- ["All Engines Table" on page 1293](#): Available metrics from the Hawk microagent for each BW Engine.
- ["All Engines Grid" on page 1296](#): Displays the main health metrics and a single trend graph per engine, summarizing the status of each BW Engine.
- ["Single Engine Summary" on page 1298](#): Detailed performance metrics and alert status for a single BW Engine.

### All Engines Heatmap

Quick view of BW5 Engines status for the selected **Filter** and **Server**. Each rectangle in the heatmap represents an engine. Rectangle size represents Max Heap Size and the color represents the most severe value in the heatmap rectangle is shown for the selected Metric. By default, the maximum **Alert Severity** is shown:

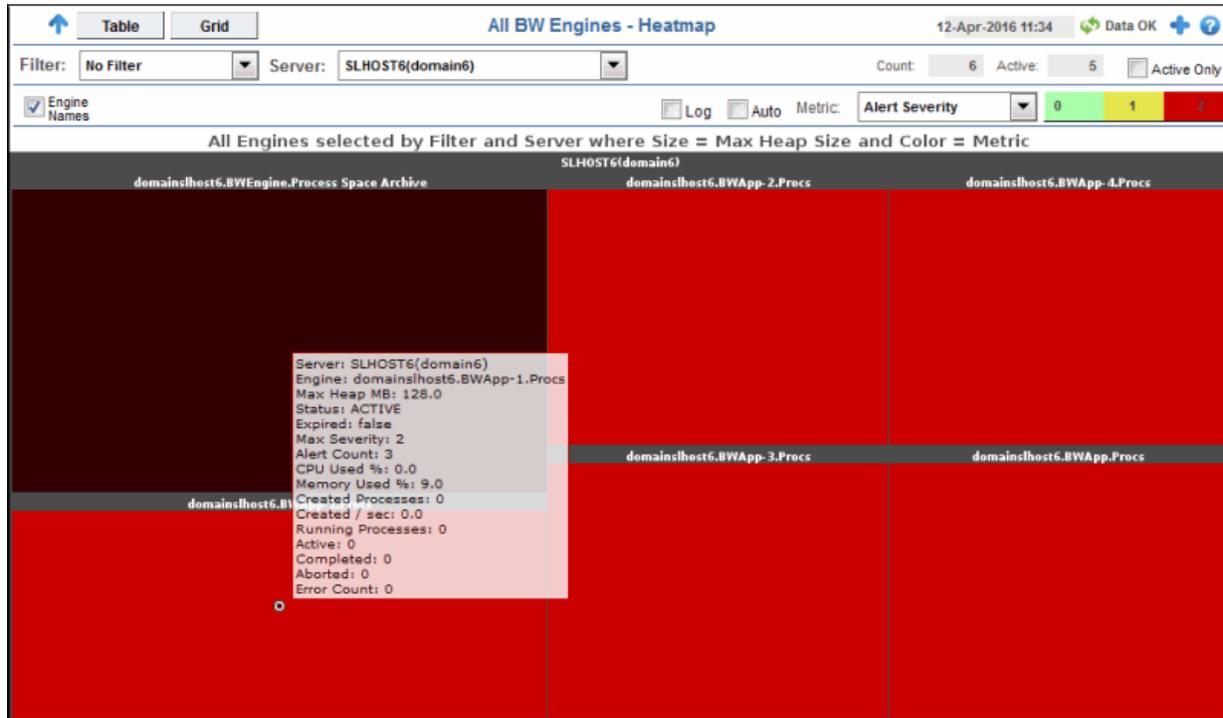
Values range from **0** - **2**, as indicated in the color gradient  bar, where **2** is the highest Alert Severity:

- Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
- Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
- Green indicates that no metrics have exceeded their alert thresholds.

Mouseover to see the following performance metrics:

- **Server**: Server agent name.
- **Engine**: Engine name.
- **Max Heap MB**: Maximum heap allocated to this engine for the JVM.
- **Status**: ACTIVE, STOPPED or LIMITED.
- **Expired**: When checked, data has not been received from this host in the specified amount of time.
- **Alert Count**: Number of current alerts
- **CPU Used %**: Percent of server CPU used by engine.
- **Memory Used %**: Percentage of allocated memory currently consumed by this engine from within the JVM. Equal to the value of:  $(100 * \text{UsedBytes}) / \text{MaxBytes}$ . NOTE: Percent used is Long.
- **Created Processes**: The total number of processes created.
- **Created / sec**: The number of processes created per second.
- **Running Processes**: The number of currently running processes.
- **Active**: The number of currently active processes.
- **Completed**: The total number of completed processes.
- **Aborted**: The total number of aborted processes.
- **Error Count**: The total number of errors.

Click on a node to drill down to the “Single Engine Summary” display to look at number of processes running, threads, history of memory utilization and other performance metrics for a specific engine. Mouse-over nodes to view details about engine performance and status.



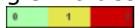
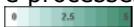
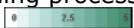
**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- The number of items currently in the display.
- Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

**Filter By:**

The display might include these filtering options:

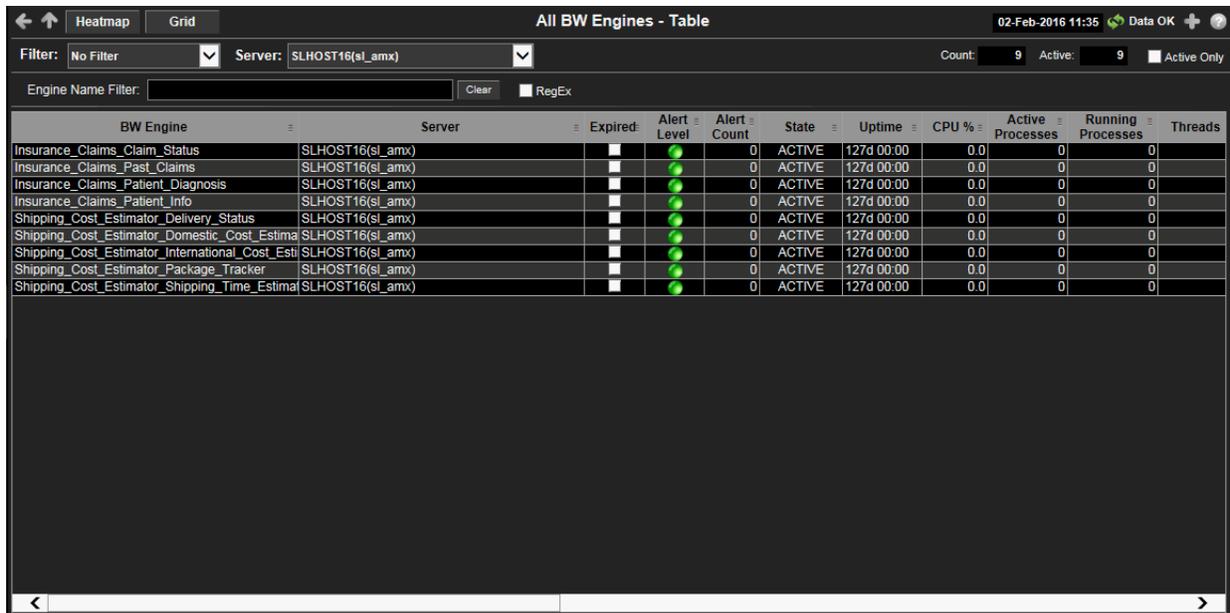
- Filter:** Choose a filter to show data for in the display. By default, the **Filter:** drop-down menu only contains the **No Filter** option. To create your own filtering options, see **Creating Customized Filters** in the User’s Guide.
- Server:** Choose a server to show data for in the display.
- Count:** The total number of engines in the display.
- Active** Number of engines currently active.
- Active Only** If selected, only engines with a status of ACTIVE are displayed. Otherwise, if deselected, all engines for the given Filter/Server selection are displayed.
- Engine Names** Select this check box to display the names of the engines above their respective rectangles in the heatmap.

<b>Log</b>	Select to enable a logarithmic scale. Use <b>Log Scale</b> to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. <b>Log Scale</b> makes data on both scales visible by applying logarithmic values rather than actual values to the data.
<b>Auto</b>	Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value. NOTE: Some metrics auto-scale automatically, even when <b>Auto</b> is not selected.
<b>Metric</b>	Choose a metric to view in the display.
<b>Alert Severity</b>	<p>The maximum level of alerts in the heatmap rectangle. Values range from <b>0</b> - <b>2</b>, as indicated in the color gradient  bar, where <b>2</b> is the highest Alert Severity:</p> <ul style="list-style-type: none"> <li> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Alert Count</b>	The total number of critical and warning alerts in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.
<b>CPU Used%</b>	The percent (%) CPU used in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Memory Used%</b>	The percent (%) memory used in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Active Processes</b>	The number of currently active processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Running Processes</b>	The number of currently running processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Created Processes</b>	The number of created processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

- Created/ sec** The number of created processes in the heatmap rectangle, per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
- Error Count** The total number of errors in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.

## All Engines Table

Each row in the table is an engine. Metrics are made available by the Hawk microagent for the engine (for details, refer to documentation for TIBCO BusinessWorks Administration, Appendix A: TIBCO Hawk Microagent Methods). Click on an row to drill down to the “Single Engine Summary” display.



BW Engine	Server	Expired	Alert Level	Alert Count	State	Uptime	CPU %	Active Processes	Running Processes	Threads
Insurance_Claims_Claim_Status	SLHOST16(sl_amx)	<input type="checkbox"/>		0	ACTIVE	127d 00:00	0.0	0	0	0
Insurance_Claims_Past_Claims	SLHOST16(sl_amx)	<input type="checkbox"/>		0	ACTIVE	127d 00:00	0.0	0	0	0
Insurance_Claims_Patient_Diagnosis	SLHOST16(sl_amx)	<input type="checkbox"/>		0	ACTIVE	127d 00:00	0.0	0	0	0
Insurance_Claims_Patient_Info	SLHOST16(sl_amx)	<input type="checkbox"/>		0	ACTIVE	127d 00:00	0.0	0	0	0
Shipping_Cost_Estimator_Delivery_Status	SLHOST16(sl_amx)	<input type="checkbox"/>		0	ACTIVE	127d 00:00	0.0	0	0	0
Shipping_Cost_Estimator_Domestic_Cost_Estima	SLHOST16(sl_amx)	<input type="checkbox"/>		0	ACTIVE	127d 00:00	0.0	0	0	0
Shipping_Cost_Estimator_International_Cost_Esti	SLHOST16(sl_amx)	<input type="checkbox"/>		0	ACTIVE	127d 00:00	0.0	0	0	0
Shipping_Cost_Estimator_Package_Tracker	SLHOST16(sl_amx)	<input type="checkbox"/>		0	ACTIVE	127d 00:00	0.0	0	0	0
Shipping_Cost_Estimator_Shipping_Time_Estima	SLHOST16(sl_amx)	<input type="checkbox"/>		0	ACTIVE	127d 00:00	0.0	0	0	0

**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** , **Table** open commonly accessed displays.
- 6,047** The number of items currently in the display.

- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

**Filter By:**

The display might include these filtering options:

- Filter:** Choose a filter to show data for in the display. By default, the **Filter:** drop-down menu only contains the **No Filter** option. To create your own filtering options, see **Creating Customized Filters** in the User’s Guide.

<b>Server:</b>	Choose a server to show data for in the display.
<b>Count</b>	Number of engines currently being displayed.
<b>Active</b>	Number of engines currently active.
<b>Active Only</b>	If selected, only engines with a status of ACTIVE are displayed. Otherwise, if deselected, all engines for the given Filter/Server selection are displayed.
<b>Engine Name Filter</b>	Enter all or part of engine name to view specific engines. NOTE: Wild card characters are supported.
<b>Clear</b>	Removes Engine Name Filter and all engines for the given Filter/Server selection are displayed.
<b>RegEx</b>	If selected, the specified Engine Name Filter will be interpreted as a full Regular Expression rather than a simple wildcard.

**Table:**

<b>BW Engine</b>	BW Engine name.
<b>Server</b>	Server agent name.
<b>Expired</b>	When checked, data has not been received from this host in the specified amount of time.
<b>Alert Level</b>	The most critical alert state for alerts in the row:  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
<b>Alert Count</b>	Number of current alerts
<b>State</b>	Engine status: ACTIVE, STOPPED, LIMITED, etc. (See <b>All Servers Grid</b> ).
<b>Uptime</b>	Uptime in milliseconds since the engine was started.
<b>CPU %</b>	Percent of server CPU used by engine.
<b>Active Processes</b>	Number of active processes calculated each update period using data returned by the Hawk method GetProcesses. <b>Note:</b> This column will display <b>NaN</b> or <b>Not Available</b> for any engine whose status is <b>STOPPED</b> .
<b>Running Processes</b>	Number of running processes.
<b>Threads</b>	Number of threads used by the engine.
<b>Memory Used%</b>	Percentage of allocated memory currently consumed by this engine from within the JVM. Equal to the value of: (100*UsedBytes) divided by MaxBytes. NOTE: Percent used is Long.
<b>Max Heap Size</b>	Maximum heap allocated to this engine for the JVM.
<b>Total Bytes</b>	Maximum heap memory this JVM has used.
<b>Used Bytes</b>	Total bytes of memory within the JVM currently used by the engine. Equal to value of: MaxBytes minus FreeBytes.
<b>Free Bytes</b>	Amount of available memory from within the JVM.
<b>Mem Usage KBytes</b>	Server memory in KB used by engine.

<b>Errors</b>	Total number of errors since the engine was started.
<b>Delta Errors</b>	Current number of new errors.
<b>Errors/sec</b>	Error rate per second.
<b>Created Processes</b>	The total number of processes that were created.
<b>Completed Processes</b>	The total number of processes that were completed.
<b>Aborted Processes</b>	The total number of processes that were aborted.
<b>Process ID</b>	Process ID of engine as recognized by the server.
<b>Micro Agent Instance</b>	Unique ID of the microagent reporting the metrics.
<b>Deployment</b>	Name of Deployment.
<b>Domain</b>	Name of Domain.
<b>BW Version</b>	The TIBCO BusinessWorks version currently in use on the server.
<b>Source</b>	Name of RTView Data Server sending this data (or localhost).
<b>Time Stamp</b>	Time of last update.
<b>Process Name</b>	Name of the BW Engine process on the server. Note: This information is not displayed in the table but is present in "raw" cache data.
<b>Host</b>	Host name of server. Note: This information is not displayed in the table but is present in "raw" cache data.
<b>Adapter Name</b>	Name of adapter. Note: This information is not displayed in the table but is present in "raw" cache data.
<b>Instance ID</b>	Instance ID name of the engine. Note: This information is not displayed in the table but is present in "raw" cache data.
<b>Version</b>	Engine project version number. Note: This information is not displayed in the table but is present in "raw" cache data.

## All Engines Grid

Displays the main health metrics and a single trend graph per engine, summarizing the status of each BW5 Engine. Click on an engine icon to drill down to the ["Single Engine Summary"](#) display.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** , **Table** open commonly accessed displays.
- 6,047** The number of items currently in the display.

- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

### Filter By:

The display might include these filtering options:

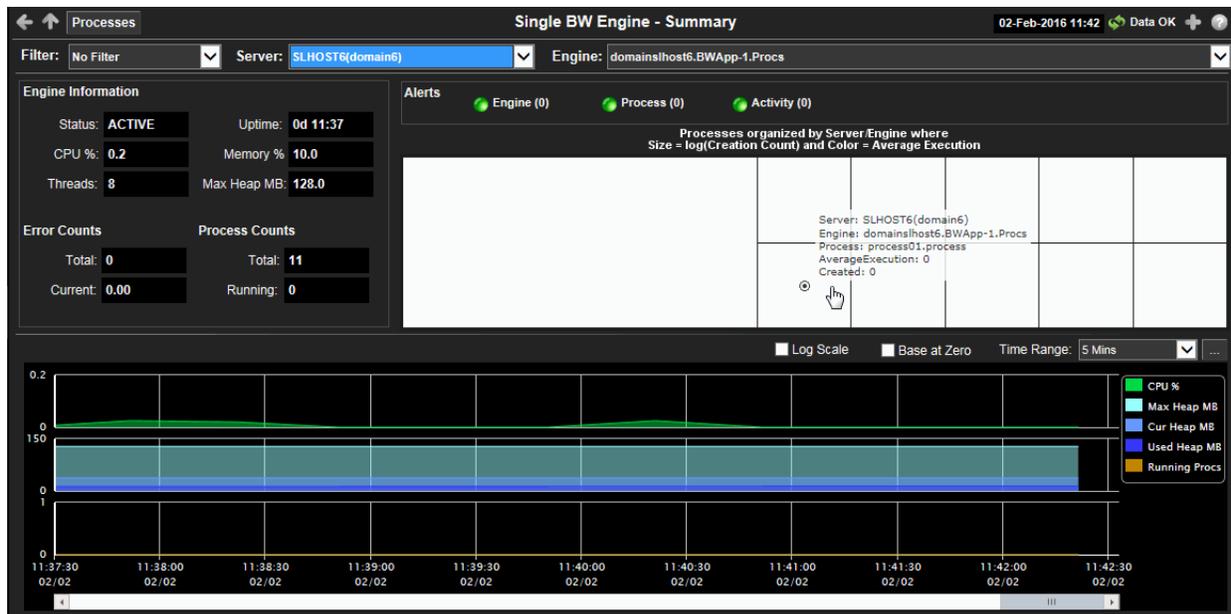
- Filter:** Choose a filter to show data for in the display. By default, the **Filter:** drop-down menu only contains the **No Filter** option. To create your own filtering options, see **Creating Customized Filters** in the User's Guide.
- Server:** Choose a server to show data for in the display.
- Count** Number of engines currently being displayed.
- Active** Number of engines currently active.
- Active Only** Toggle this setting to display active servers or all servers.
- Time Range** Choose a time range. Also sets range for the **Single Engine Summary** display. Options are:  
**All Data, 2 Mins, 5 Mins, 20 Mins, 1 Hour, 2 Hours, 4 Hours, 8 Hours, 24 Hours, 2 Days and 7 Days.**

## Fields and Data

<b>Engine Name</b>	Name of the engine.
<b>Status</b>	Indicates the current state of the engine: <ul style="list-style-type: none"> <li>• <b>ACTIVE</b> Indicates the BW microagent is providing live data and the engine is assumed active.</li> <li>• <b>SUSPENDED</b> This state is reported by the BW microagent.</li> <li>• <b>STANDBY</b> This state is reported by the BW microagent.</li> <li>• <b>STOPPING</b> This state is reported by the BW microagent.</li> <li>• <b>STOPPED</b> This state is reported by the BW microagent.</li> <li>• <b>LIMITED</b> Live data has been received from TIBCO, but deployment data from the custom RTView microagent has not been received.</li> <li>• <b>EXPIRED</b> Indicates the server associated with the engine is unavailable or stopped sending data.</li> </ul>
<b>CPU Usage%</b>	Percent of server CPU in use.
<b>Mem Usage%</b>	Available physical memory (MB) remaining.
<b>Error Rate</b>	Number of errors accumulated per second.
<b>Total Processes</b>	Number of process definitions for this engine.
<b>Active Processes</b>	Number of process instances currently active.
<b>Trend Graphs</b>	Traces data for the server.
<b>CPU</b>	Traces percent of server CPU in use.
<b>MEM</b>	Traces available physical memory remaining.
<b>PROCS</b>	Traces total number of active processes.

## Single Engine Summary

Several views show historical and current performance metrics for a single engine, including the number of processes running, threads, history of memory utilization, and trend graphs of memory utilization. In this display, when an engine is **Stopped** the engine name is appended with **(X)**, the background color is light red and Uptime is zero.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** , **Table** open commonly accessed displays.
- 6,047** The number of items currently in the display.

- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

**Note:** Clicking the **Processes** button in the Title Bar takes you to the **"All Processes Heatmap"**. Clicking the **JVM** button, which is automatically enabled when a JMX connection is defined for the engine, takes you to the **JVM CPU/Mem Summary** display. See **Enable Monitoring Via JMX** for more information on enabling a JMX connection.

### Filter By:

The display might include these filtering options:

- Filter:** Choose a filter to show data for in the display. By default, the **Filter:** drop-down menu only contains the **No Filter** option. To create your own filtering options, see **Creating Customized Filters** in the User's Guide.
- Server:** Choose a server to show data for in the display.
- Engine:** Choose an engine to show data for in the display. An engine is not running when the engine name is appended with **(X)**.

## Fields and Data

### Engine Information

<b>Status</b>	<b>ACTIVE</b>	The BW microagent is providing live data and the engine is assumed active.
	<b>SUSPENDED</b>	This state is reported by the BW microagent.
	<b>STANDBY</b>	This state is reported by the BW microagent.
	<b>STOPPING</b>	This state is reported by the BW microagent.
	<b>STOPPED</b>	This state is reported by the BW microagent.
	<b>LIMITED</b>	Live data has been received from TIBCO, but deployment data from the custom RTView MicroAgent has not been received.
	<b>EXPIRED</b>	The associated server for the engine is currently in an EXPIRED state and is unavailable or stopped sending data.
<b>Uptime</b>		Days hours and minutes since the engine was started.
<b>CPU%</b>		Percent of server CPU used by engine.
<b>Memory %</b>		Available physical memory remaining (in MB).
<b>Threads</b>		Number of threads used by this engine
<b>Max Heap MB</b>		Maximum heap allocated to this engine for the JVM.

### Error Counts

<b>Total</b>	Total errors accumulated by this engine.
<b>Current</b>	Number of errors accumulated this update cycle.

### Process Counts

<b>Total</b>	A BW Engine runs processes by creating instances of process definitions and making them active. A given process instance has a lifetime during which it may be suspended, swapped, queued, etc. until it is either completed or aborted. The Total value is calculated using the Hawk method named GetProcessDefinitions that returns statistics about the instances of each process definition including cumulative counts of completed, aborted, suspended, etc.
<b>Running</b>	Total number of running process instances. This number is calculated using the Hawk method named GetProcessCount. It is displayed in the Monitor Engines Table as RunningProcesses. The trend below displays the same value over time as Running Procs.

### Alerts

Indicates the greatest severity level and the number of open **Engine**, **Process**, and **Activity** alerts for the selected engine. Values range from **0** to **2**, where **2** is the greatest Severity:

- One or more alerts exceeded their ALARM LEVEL threshold.
- One or more alerts exceeded their WARNING LEVEL threshold.
- No alert thresholds have been exceeded.

Click on the alert indicator to display the **BW Engine - Tables** display, which contains a table listing the current alerts for the selected engine.

The screenshot shows the 'BW Engine - Tables' window with the following data:

BW Engine Table									
State	CPU %	Running Processes	Threads	Memory Used %	Max Heap Size	Total Bytes	Used Bytes	Free Bytes	Mem Used
LIMITED	0.0	0	8	0.0	0	0	0	0	0

Process Totals Table						
Aborted	Active	AverageElapsed	AverageExecution	Checkpointed	Completed	Created
1107306	0	0.0	0.0	0	664628	1771934

Current Alerts for Selected BW Engine			
Time	Alert Name	Alert Index	Alert Text

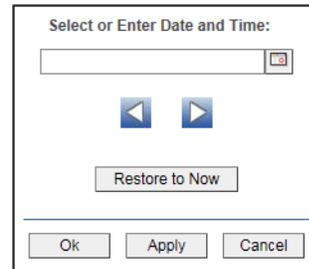
- Engine** Number of engine alerts and the most critical alert state for the engine:
- Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
  - Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
  - Green indicates that no metrics have exceeded their alert thresholds.
- Process** Number of process alerts and the most critical alert state for the engine:
- Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
  - Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
  - Green indicates that no metrics have exceeded their alert thresholds.
- Activity** Number of activity alerts and the most critical alert state for the engine:
- Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
  - Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
  - Green indicates that no metrics have exceeded their alert thresholds.

### Heatmap

Shows processes organized by Server/Engine where Size = Creation Count and Color = Average Execution. Click on a node to drill down to a specific engine.

### Trend Graphs

- Log Scale** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Base at Zero** Select to use zero (0) as the Y axis minimum for all graph traces.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## BW5 Processes

These displays present performance metrics for BW5 processes. Displays in this View are:

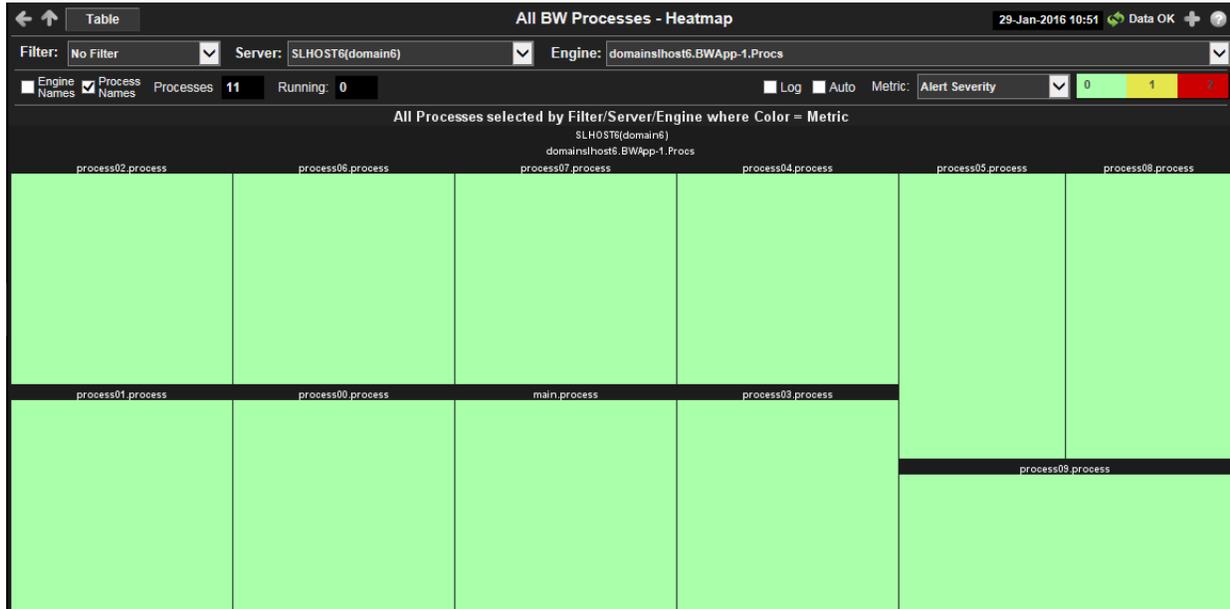
- [“All Processes Heatmap” on page 1301](#): Displays process execution metrics for all BW Engines.
- [“All Processes Table” on page 1305](#): Each row in the table displays all available metrics from the Hawk microagent for a process.
- [“Single Process Summary” on page 1308](#): Several views show historical and current metrics for a single process, including average execution times and execution counts.

## All Processes Heatmap

Summary view of processes can show the execution times for all processes on all engines or you can filter to look at specific servers or engines. Each rectangle (node) in the heatmap represents a process. Move your mouse over a node to display current metrics. Click on a node to drill-down to the [“Single Process Summary”](#) display to view specific metrics about process behavior over a specified period of time and determine which activity in the process may be causing the bottleneck.

An engine is not running when the engine name is appended with **(X)**.

Mouse-over any node to display the current values for the metric selected from the **Metric** drop-down menu.



#### Title Bar (possible features are):

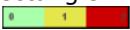
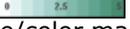
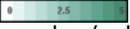
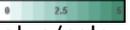
- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- The number of items currently in the display.

- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

#### Filter By:

The display might include these filtering options:

- Filter:** Choose a filter to show data for in the display. By default, the **Filter:** drop-down menu only contains the **No Filter** option. To create your own filtering options, see **Creating Customized Filters** in the User's Guide.
- Server:** Choose a server to show data for in the display.
- Engine:** Choose an engine to show data for in the display. An engine is not running when the engine name is appended with **(X)**.
- Engine Names** Select this check box to display the names of the engines above their respective rectangles in the heatmap.
- Process Names** Select this check box to display the names of the processes above their respective rectangles in the heatmap.
- Processes** The total number of processes in the display.
- Running** Number of processes currently running.

<b>Log</b>	Select to enable a logarithmic scale. Use <b>Log Scale</b> to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. <b>Log Scale</b> makes data on both scales visible by applying logarithmic values rather than actual values to the data.
<b>Auto</b>	Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value. NOTE: Some metrics auto-scale automatically, even when <b>Auto</b> is not selected.
<b>Metric</b>	Choose a metric to view in the display.
<b>Alert Severity</b>	<p>The maximum level of alerts in the heatmap rectangle. Values range from <b>0</b> - <b>2</b>, as indicated in the color gradient  bar, where <b>2</b> is the highest Alert Severity:</p> <ul style="list-style-type: none"> <li> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Alert Count</b>	The total number of critical and warning alerts in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.
<b>Completed Count</b>	The total number of completed processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Active Count</b>	The total number of active processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Aborted Count</b>	The total number of aborted processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Suspended Count</b>	The total number of suspended processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Exec Time / sec</b>	The number of processes executed per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Created / sec</b>	The number of processes created per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

<b>Aborted / sec</b>	The number of aborted processes per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Suspended / sec</b>	The number of suspended processes per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Most Recent Exec Time</b>	The execution time for the most recently executed process in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Average Exec Time</b>	The average execution time for all processes in the heatmap rectangle, calculated by dividing the delta execution time for the interval by the delta completed, or the number of process instances that completed in the interval. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Most Recent Elapsed Time</b>	The elapsed time for the most recently executed process in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Average Elapsed Time</b>	The average elapsed time for all processes in the heatmap rectangle, calculated by dividing the delta elapsed time for the interval by the delta completed, or the number of process instances that completed in the interval. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

## All Processes Table

Select a server and engine from the drop-down menus. Each row in the table is a different engine. The table displays all metrics available from the Hawk microagent for an engine. (Refer to documentation for TIBCO BusinessWorks Administration, see Appendix A: TIBCO Hawk Microagent Methods).

Click on a row in the table to drill down to the “Single Engine Summary” display.

BW Engine	Server	BW Process	Expired	Alert Level	Alert Count	Time Since Last Update	Total Exec Time	Exe
domainshost6.BWApp-1.Procs	SLHOST6(domain6)	main_process		Red	2	0	0	
domainshost6.BWApp-1.Procs	SLHOST6(domain6)	process00_process		Green	0	0	0	
domainshost6.BWApp-1.Procs	SLHOST6(domain6)	process01_process		Green	0	0	0	
domainshost6.BWApp-1.Procs	SLHOST6(domain6)	process02_process		Green	0	0	0	
domainshost6.BWApp-1.Procs	SLHOST6(domain6)	process03_process		Green	0	0	0	
domainshost6.BWApp-1.Procs	SLHOST6(domain6)	process04_process		Green	0	0	0	
domainshost6.BWApp-1.Procs	SLHOST6(domain6)	process05_process		Green	0	0	0	
domainshost6.BWApp-1.Procs	SLHOST6(domain6)	process06_process		Green	0	0	0	
domainshost6.BWApp-1.Procs	SLHOST6(domain6)	process07_process		Green	0	0	0	
domainshost6.BWApp-1.Procs	SLHOST6(domain6)	process08_process		Green	0	0	0	
domainshost6.BWApp-1.Procs	SLHOST6(domain6)	process09_process		Green	0	0	0	

**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** **Table** open commonly accessed displays.
- 6,047** The number of items currently in the display.
- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

### Filter By:

The display might include these filtering options:

- Filter:** Choose a filter to show data for in the display. By default, the **Filter:** drop-down menu only contains the **No Filter** option. To create your own filtering options, see **Create Customized Filters** for more information.
- Server:** Choose a server to show data for in the display.
- Engine:** Choose an engine to show data for in the display. An engine is not running when the engine name is appended with **(X)**.

### Table:

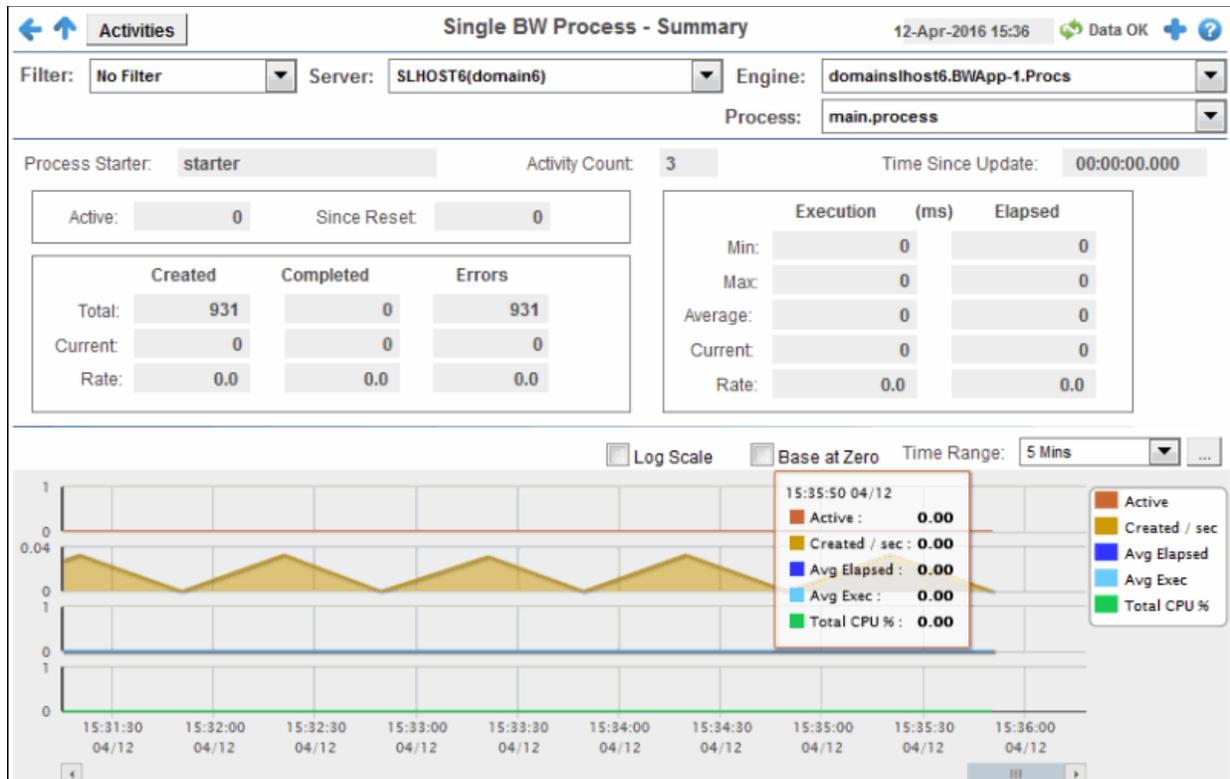
- BW Engine** BW Engine name.
- Server** Server agent name.

<b>BW Process</b>	The name of the process.
<b>Expired</b>	When checked, data has not been received from this host in the specified amount of time.
<b>Alert Level</b>	The most critical alert state for alerts in the row:  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
<b>Alert Count</b>	Number of current alerts
<b>Active</b>	Number of active processes.
<b>Total CPU</b>	Total CPU usage in percent.
<b>Created/sec</b>	Change in Created per second.
<b>Completed/sec</b>	Change in Completed per second.
<b>Delta Created</b>	Change in Created this update.
<b>Delta Completed</b>	Change in Completed this update.
<b>Created</b>	Number of process instances created for this process definition.
<b>Completed</b>	Number of process instances successfully completed.
<b>Total Exec Time</b>	Total execution time (in milliseconds) for all successfully completed process instances.
<b>Delta Exec Time</b>	Execution time accumulated during the current polling period.
<b>Exec Time/sec</b>	Delta execution time per second.
<b>Min Exec Time</b>	Execution time (in milliseconds) of the process instance that has completed in the shortest amount of execution time.
<b>Max Exec Time</b>	Execution time (in milliseconds) of the process instance that has completed in the longest amount of execution time.
<b>Average Exec Time</b>	Average execution time (in milliseconds) for all successfully completed process instances.
<b>Recent Exec Time</b>	Execution time (in milliseconds) of the most recently completed process instance.
<b>Total Elapsed Time</b>	Total elapsed time (in milliseconds) for all successfully completed process instances.
<b>Delta Elapsed Time</b>	Elapsed time accumulated during the current polling period.
<b>Elapsed Time/sec</b>	Delta elapsed time per second.
<b>Min Elapsed Time</b>	Elapsed clock time (in milliseconds) of the process instance that has completed in the shortest amount of elapsed time.
<b>Max Elapsed Time</b>	Elapsed clock time (in milliseconds) of the process instance that has completed in the longest amount of elapsed time.
<b>Average Elapsed Time</b>	Average elapsed clock time (in milliseconds) for all successfully completed process instances.

<b>Recent Elapsed Time</b>	Elapsed clock time (in milliseconds) of the most recently completed process instance.
<b>Aborted</b>	Number of times process instances have been aborted.
<b>Delta Aborted</b>	Change in Aborted this update.
<b>Aborted/sec</b>	Change in Aborted per second.
<b>Queued</b>	Number of times process instances have been queued for execution.
<b>Delta Queued</b>	Change in Queued this update.
<b>Queued/sec</b>	Change in Queued per second.
<b>Suspended</b>	Number of times process instances have been suspended.
<b>Delta Suspended</b>	Change in Suspended this update.
<b>Suspended/sec</b>	Change in Suspended per second.
<b>Checkpointed</b>	Number of times process instances have executed a checkpoint.
<b>Delta Checkpointed</b>	Change in Checkpointed this update.
<b>Checkpointed/sec</b>	Change in Checkpointed per second.
<b>Swapped</b>	Number of times process instances have been swapped to disk.
<b>Delta Swapped</b>	Change in Swapped this update.
<b>Swapped/sec</b>	Change in Swapped per second.
<b>Time Since Last Update</b>	Time since the last update.
<b>Domain</b>	Name of TIBCO Domain.
<b>Starter</b>	Name of the process starter for the process.
<b>MicroAgent Instance</b>	Unique ID of the microagent reporting the metrics.
<b>CountSince Reset</b>	Number of process instances that have completed since the last reset of the statistics.
<b>Source</b>	Name of RTView Data Server sending this data (or localhost).
<b>Time Stamp</b>	Time of last update.

## Single Process Summary

Detailed performance metrics and alert status for a single BW process. Select a server, engine and process from the drop-down menus. The background color of the display is red when the selected engine is stopped.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** , **Table** open commonly accessed displays.
- 6,047** The number of items currently in the display.

- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

### Filter By:

The display might include these filtering options:

- Filter:** Choose a filter to show data for in the display. By default, the **Filter:** drop-down menu only contains the **No Filter** option. To create your own filtering options, see **Creating Customized Filters** in the User's Guide.
- Server:** Choose a server to see metrics for.
- Engine:** Choose a server to see metrics for. An engine is not running when the engine name is appended with **(X)**.
- Process:** Choose a process to see metrics for.

<b>Process Starter</b>	Name of the process starter for the process.
<b>Activity Count</b>	Number of activities defined for this process.
<b>Time Since Update</b>	Time since the last update to file of statistics.
<b>Active</b>	Number of active instances for this process definition. This number is calculated using the Hawk method named GetProcesses. This method returns information about process instances that are active at the time of update. The value here displays the current total count of all active instances discovered for this process definition. The trend below displays the same value over time.
<b>Since Reset</b>	Number of activity executions that have completed since the last reset of the statistics. This is the number retrieved from the Hawk method named GetProcessDefinition which returns ExecutionCountSinceReset.

### Execution Counts

Most recent execution counts for this process.

<b>Created</b>	<b>Total</b>	Number of process instances created for this process definition.
	<b>Current</b>	Number of process instances created this update cycle.
	<b>Rate</b>	Number of process instances created per second.
<b>Completed</b>	<b>Total</b>	Number of process instances successfully completed.
	<b>Current</b>	Number of process instances successfully completed this update cycle.
	<b>Rate</b>	Number of process instances successfully completed per second.
<b>Errors</b>	<b>Total</b>	Number of errors accumulated by all process instances.
	<b>Current</b>	Number of errors accumulated this update cycle.
	<b>Rate</b>	Number of errors accumulated per second.

### Execution (ms) Elapsed

Execution and elapsed times in milliseconds for this process.

<b>Min</b>	Shortest time of any process instance.
<b>Max</b>	Longest time of any process instance.
<b>Average</b>	Average time across all successfully completed process instances.
<b>Current</b>	Time accumulated this update cycle.
<b>Rate</b>	Time accumulated per second.

### Trend Graphs

- **Active:** Traces the number of currently active processes.
- **Created / sec:** Traces the number of created processes per second.
- **Avg Elapsed:** Traces the average number of elapsed processes. This value is calculated by dividing the delta elapsed time for the interval by the delta completed, or the number of process instances that completed in the interval.
- **Avg Exec:** Traces the average number of executed processes. This value is calculated by dividing the delta executed time for the interval by the delta completed, or the number of process instances that completed in the interval.
- **Total CPU %:** Traces CPU utilization by processes, in percent.

- Log Scale** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Base at Zero** Select to use zero (**0**) as the Y axis minimum for all graph traces.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## BW5 Activities

These displays present performance metrics for BW5 activities. Displays in this View are:

- ["All Activities Heatmap" on page 1310](#): Displays execution performance metrics for all BW activities.
- ["All Activities Table" on page 1313](#): Each row in the table displays all available metrics from the Hawk microagent for an activity.
- ["Single Activity Summary" on page 1316](#): Historical and current performance metrics for a single activity, including average execution times and execution counts.

## All Activities Heatmap

Summary view of activities shows the execution times for all activities on all engines, or you can filter to look at specific servers, engines or processes. An engine is not running when the engine name is appended with **(X)**.

Move your mouse over a node to display current metrics. Click on a node to drill down to the “Single Activity Summary” display to view specific metrics about activity behavior over a specified period of time.



**Title Bar** (possible features are):

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu, Table open commonly accessed displays.
- 6,047 The number of items currently in the display.
- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

**Filter By:**

The display might include these filtering options:

- Filter:** Choose a filter to show data for in the display. By default, the **Filter:** drop-down menu only contains the **No Filter** option. To create your own filtering options, see **Creating Customized Filters** in the User’s Guide.
- Server:** Choose a server to show data for in the display.
- Engine:** Choose an engine to show data for in the display. An engine is not running when the engine name is appended with **(X)**.
- Process** Select from the menu to view activities running on a specific process or all processes.
- OK** Number of activities that reported their Last Return Code as **OK**.
- Error** Number of activities that reported their Last Return Code as **Error**.
- Dead** Number of activities that reported their Last Return Code as **Dead**.
- Engine Names** Select this check box to display the names of the engines above their respective rectangles in the heatmap.

<b>Process Names</b>	Select this check box to display the names of the processes above their respective rectangles in the heatmap.
<b>Activity Names</b>	Select this check box to display the names of the activities above their respective rectangles in the heatmap.
<b>Log</b>	Select to enable a logarithmic scale. Use <b>Log Scale</b> to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. <b>Log Scale</b> makes data on both scales visible by applying logarithmic values rather than actual values to the data.
<b>Auto</b>	Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value. NOTE: Some metrics auto-scale automatically, even when <b>Auto</b> is not selected.
<b>Metric</b>	Choose a metric to view in the display.
<b>Alert Severity</b>	<p>The maximum level of alerts in the heatmap rectangle. Values range from <b>0</b> - <b>2</b>, as indicated in the color gradient  bar, where <b>2</b> is the highest Alert Severity:</p> <ul style="list-style-type: none"> <li> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Alert Count</b>	The total number of critical and warning alerts in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.
<b>Exec Count</b>	The total number of executed processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Error Count</b>	The total number of errors in the heatmap rectangle. The color gradient  bar populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Exec Time / sec</b>	The number of processes executed per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Errors / sec</b>	The number of errors per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

**Most Recent Exec Time**

The execution time for the most recently executed process in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

**Max Exec Time**

The maximum execution time for executed processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

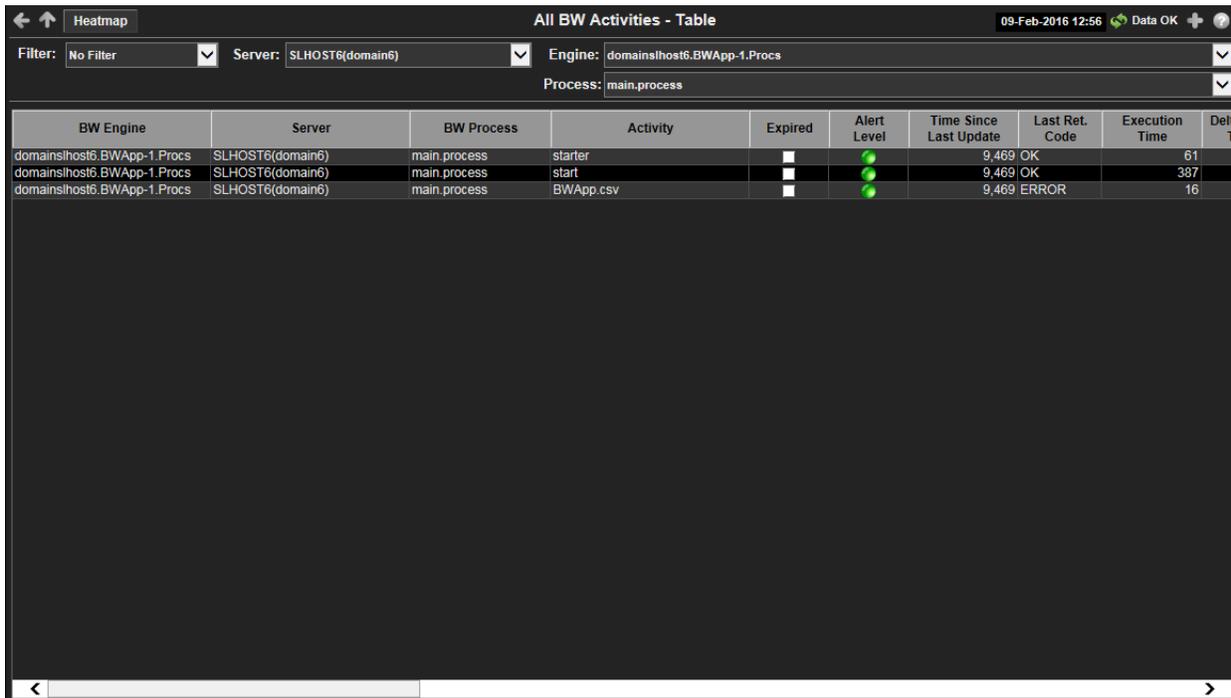
**All Activities Table**

Select a server, engine and process from the drop-down menus. Each row in the table displays all metrics available from the Hawk microagent for an activity. (Refer to documentation for TIBCO BusinessWorks Administration, see Appendix A: TIBCO Hawk Microagent Methods).

Click on a row in the table to drill down to the “Single Activity Summary” display to view specific metrics about activity behavior over a specified period of time.

When the background/foreground color of a row changes color, the associated engine for the activity is currently in an EXPIRED state. An engine is EXPIRED when the associated server is unavailable or stopped sending data.

An EXPIRED activity and the associated engine are deleted from displays when the associated server exceeds its specified threshold. Processes associated with the engine are also deleted from displays.



BW Engine	Server	BW Process	Activity	Expired	Alert Level	Time Since Last Update	Last Ret. Code	Execution Time	Del
domainshost6.BWApp-1.Procs	SLHOST6(domain6)	main.process	starter	<input type="checkbox"/>		9,469 OK		61	
domainshost6.BWApp-1.Procs	SLHOST6(domain6)	main.process	start	<input type="checkbox"/>		9,469 OK		387	
domainshost6.BWApp-1.Procs	SLHOST6(domain6)	main.process	BWApp.csv	<input type="checkbox"/>		9,469 ERROR		16	

**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- The number of items currently in the display.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

Open the **Alert Views - RTView Alerts Table** display.

**Filter By:**

The display might include these filtering options:

<b>Filter:</b>	Choose a filter to show data for in the display. By default, the <b>Filter:</b> drop-down menu only contains the <b>No Filter</b> option. To create your own filtering options, see <b>Creating Customized Filters</b> in the User's Guide.
<b>Server:</b>	Choose a server to show data for in the display.
<b>Engine:</b>	Select from the menu to view activities running on a specific engine or all engines. An engine is not running when the engine name is appended with <b>(X)</b> .
<b>Process:</b>	Select from the menu to view activities running on a specific process or all processes.

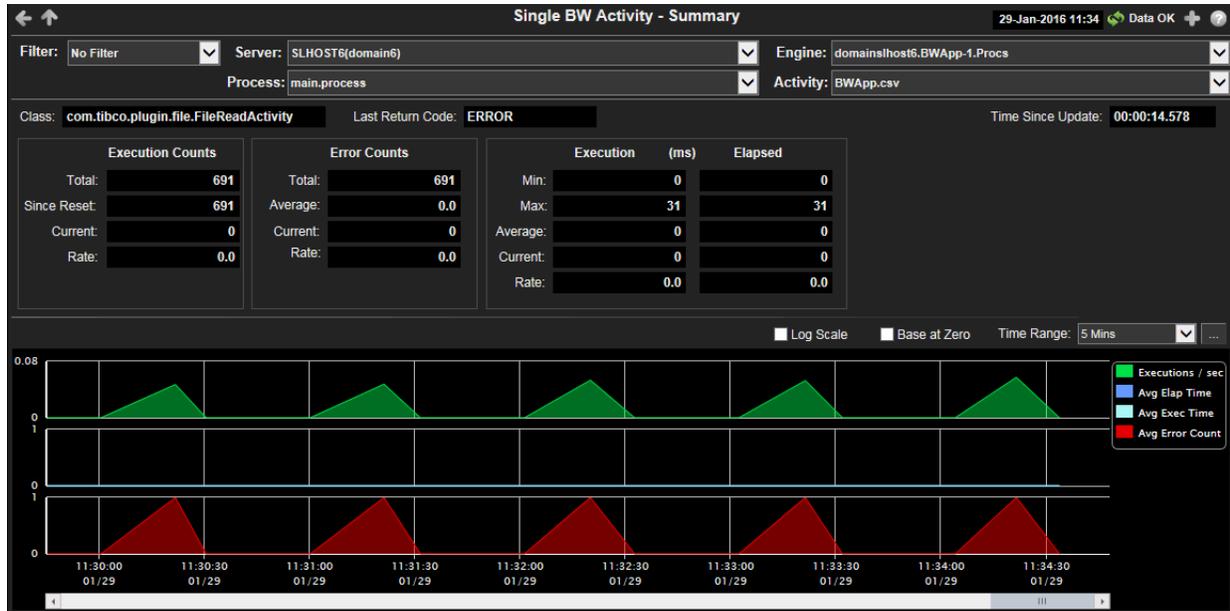
**Table:**

<b>BW Engine</b>	Name of BW Engine.
<b>Server</b>	Name of Server agent.
<b>BW Process</b>	Name of the BW engine Process on the Server.
<b>Activity</b>	Name of activity.
<b>Expired</b>	When checked, data has not been received from this host in the specified amount of time.
<b>Alert Level</b>	The most critical alert state for alerts in the row: <ul style="list-style-type: none"> <li> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Time Since Last Update</b>	Time since the last update.
<b>Last Ret(urn) Code</b>	Status code (OK DEAD ERROR) returned by most recent execution of this activity.
<b>Execution Time</b>	Time (in milliseconds) used by all executions of this activity. NOTE: This does not include wait time for Sleep, Call Process, and Wait For... activities.
<b>Delta Exec(ution) Time</b>	Execution time accumulated this update cycle.
<b>Exec(ution) Time / sec</b>	Execution time accumulated per second.

<b>Min Exec(ution) Time</b>	Time (in milliseconds) of the activity that has the shortest execution time.
<b>Max Exec(ution) Time</b>	Time (in milliseconds) of the activity that has the longest execution time.
<b>Elapsed Time</b>	Elapsed clock time (in milliseconds) used by all executions of this activity. NOTE: This does not include wait time for Sleep, Call Process, and Wait For... activities.
<b>Delta Elapsed Time</b>	Change in ElapsedTime this update.
<b>Elapsed Time/sec</b>	Change in ElapsedTime per second.
<b>Min Elapsed Time</b>	Elapsed clock time (in milliseconds) of the activity that has the shortest execution time.
<b>Max Elapsed Time</b>	Elapsed clock time (in milliseconds) of the activity that has the longest execution time.
<b>Executions</b>	Number of times the activity has been executed.
<b>Delta Exec(ution)</b>	Change in ExecutionCount this update.
<b>Executions/sec</b>	Change in ExecutionCount per second.
<b>Errors</b>	Total number of executions of the activity that have returned an error.
<b>Delta Errors</b>	Change in ErrorCount this update.
<b>Errors/sec</b>	Change in ErrorCount per second.
<b>Domain</b>	Name of TIBCO Domain.
<b>ActivityClass</b>	Name of the class that implements the activity.
<b>CalledProcessDefs</b>	A comma-separated list of definitions called by this activity.
<b>Tracing</b>	<ul style="list-style-type: none"> <li>• <b>true</b> Tracing is enabled for this activity.</li> <li>• <b>false</b> Tracing is disabled for this activity.</li> </ul>
<b>MicroAgentInstance</b>	Unique ID of the microagent reporting the metrics.
<b>ExecutionCountSince Reset</b>	Number of times the activity has been executed since the last reset of the statistics.
<b>Source</b>	Name of RTView Data Server sending this data (or localhost).
<b>Time Stamp</b>	Time of this update.

## Single Activity Summary

Detailed performance metrics and alert status for a single BW activity. In this display, when an engine associated with the activity is **Stopped** the engine name is appended with **(X)** and the background color is light red.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- 6,047** The number of items currently in the display.

- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

### Filter By:

The display might include these filtering options:

- Filter:** Choose a filter to show data for in the display. By default, the **Filter:** drop-down menu only contains the **No Filter** option. To create your own filtering options, see **Creating Customized Filters** in the User's Guide.
- Server:** Select from the menu to view processes running on a specific server.
- Engine:** Select from the menu to view processes running on a specific engine. An engine is not running when the engine name is appended with **(X)**.
- Process:** Select from the menu to view summary details for a specific process.
- Activity** Select from the menu to view summary details for a specific activity.

**Class** Name of the activity class.

**Last Return Code** Last return code reported from this activity.

**Time Since Update** Time since the last update.

#### Execution Counts

Most recent execution counts for this activity.

**Total** Number of times the activity has been executed.

**Since Reset** Number of times the activity has been executed since the last Hawk reset of the statistics.

**Current** Change in ExecutionCount this update.

**Rate** Change in Execution Count per second.

#### Error Counts

Most recent error counts for this activity.

**Total** Number of errors accumulated by all activities.

**Average** Average number of errors accumulated by all activities.

**Current** Number of errors accumulated this update cycle.

**Rate** Number of errors accumulated per second.

#### Execution (ms) Elapsed

Execution and elapsed times in milliseconds for this activity.

**Min** Shortest time of any activity instance.

**Max** Longest time of any activity instance.

**Average** Average time across all successfully completed activity instance.

**Current** Time accumulated this update cycle.

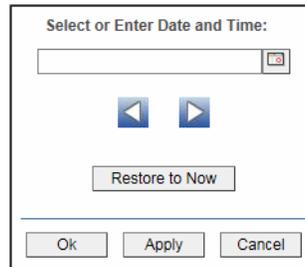
**Rate** Time accumulated per second.

#### Trend Graphs

**Log Scale** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.

**Base at Zero** Select to use zero (**0**) as the Y axis minimum for all graph traces.

**Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## BW5 Servers

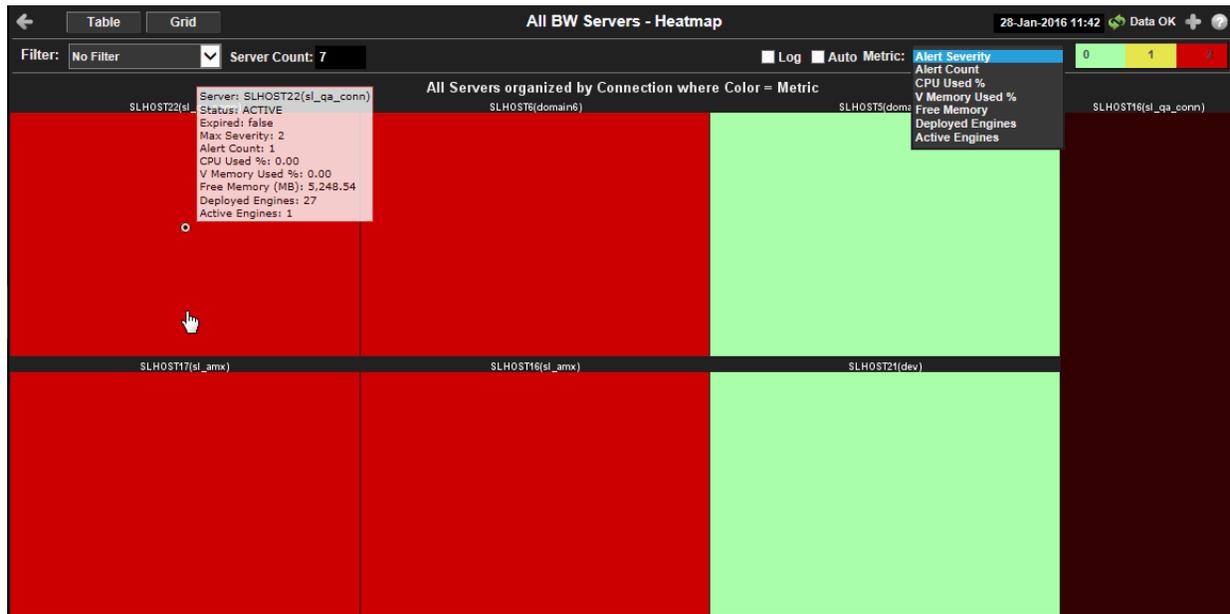
These displays present performance data for your BusinessWorks system. Displays in this View are:

- ["All Servers Heatmap" on page 1318](#)
- ["All Servers Table" on page 1321](#)
- ["All Servers Grid" on page 1322](#)
- ["Single Server Summary" on page 1324](#)
- ["Server Processes" on page 1326](#)
- ["Single Server Process - Summary" on page 1327](#)

## All Servers Heatmap

Quick view of BW Servers status determined by selected Filter, organized by Connection (host) and where color equals the selected Metric. Each rectangle (node) in the heatmap represents a server.

Click on a node to drill down to the “Single Server Summary” display and view metrics for a particular server. Mouse-over any node to display the current values for the metric selected from the Metric drop-down menu.



**Title Bar (possible features are):**

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu, Table open commonly accessed displays.
- 6,047 The number of items currently in the display.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

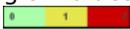
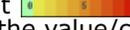
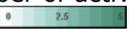
23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

Open the **Alert Views - RTView Alerts Table** display.

### Filter By:

The display might include these filtering options:

- Filter:** Choose a filter to limit data shown in the display. By default, the **Filter:** drop-down menu only contains the **No Filter** option. To create your own filtering options, see **Creating Customized Filters** in the User’s Guide.
- Server Count:** The total number of servers in the display.
- Log** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Auto** Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar’s maximum range displays the highest value. NOTE: Some metrics auto-scale automatically, even when **Auto** is not selected.
- Metric** Choose a metric to view in the display.

<b>Alert Severity</b>	<p>The maximum level of alerts in the heatmap rectangle. Values range from <b>0</b> - <b>2</b>, as indicated in the color gradient  bar, where <b>2</b> is the highest Alert Severity:</p> <ul style="list-style-type: none"> <li> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Alert Count</b>	<p>The total number of critical and warning alerts in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.</p>
<b>CPU Used%</b>	<p>The percent (%) CPU used in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.</p>
<b>V(irtual) Memory Used%</b>	<p>The percent (%) virtual memory used in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.</p>
<b>Free Memory</b>	<p>The amount of free memory in the heatmap rectangle, in megabytes. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.</p>
<b>Deployed Engines</b>	<p>The number of deployed engines in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.</p>
<b>Active Engines</b>	<p>The number of active engines in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.</p>

## All Servers Table

This table provides a list view of utilization metrics for all BW servers (represented in the All Servers Heatmap). Each row in the table contains data for a particular server. Click a column header to sort column data in numerical or alphabetical order. Click on a table row to drill down to the “Single Server Summary” display and view metrics for that particular server.

Server	Expired	Alert Level	State	CPU Usage (%)	Free Memory (MB)	V. Memory Usage (%)	BW Version	Deployed Engines	Active Engines	Source	Time Stamp
SLHOST16(sl_amx)	<input type="checkbox"/>		ACTIVE	5.95	926.28	18.97		9	9	localhost	01/28/16 11:48:30
SLHOST16(sl_qa_conn)	<input checked="" type="checkbox"/>		EXPIRED	10.74	816.28	19.91	v5.10	0	0	localhost	01/28/16 11:30:04
SLHOST17(sl_amx)	<input type="checkbox"/>		ACTIVE	0.69	3,323.74	2.20		9	9	localhost	01/28/16 11:48:21
SLHOST21(dev)	<input type="checkbox"/>		ACTIVE	4.00	2,446.26	20.80		0	0	localhost	01/28/16 11:48:49
SLHOST22(sl_qa_conn)	<input type="checkbox"/>		ACTIVE	0.00	5,249.51	0.00	v5.10	27	1	localhost	01/28/16 11:48:31
SLHOST5(domain5)	<input type="checkbox"/>		ACTIVE	17.33	1,763.04	0.71	v5.7	5	0	localhost	01/28/16 11:48:29
SLHOST6(domain6)	<input type="checkbox"/>		ACTIVE	3.52	915.39	1.68	v5.7	6	5	localhost	01/28/16 11:48:21

**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** , **Table** open commonly accessed displays.
- 6,047** The number of items currently in the display.

- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

**Filter By:**

The display might include these filtering options:

**Filter:** Choose a filter to show data for in the display. By default, the **Filter:** drop-down menu only contains the **No Filter** option. To create your own filtering options, see **Creating Customized Filters** in the User’s Guide.

**Table:**

**Server** Name of Server Agent.

**Expired** When checked, data has not been received from this host in the specified amount of time.

<b>Alert Level</b>	The most critical alert state for alerts in the row: <ul style="list-style-type: none"> <li><span style="color: red;">●</span> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li><span style="color: yellow;">●</span> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li><span style="color: green;">●</span> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>State</b>	The current status of the application. Valid values are <b>Running</b> and <b>Stopped</b> .
<b>CPU Usage (%)</b>	Percent of server CPU in use.
<b>Free Memory (MB)</b>	Available physical memory (MB) remaining.
<b>V. Memory Usage (%)</b>	Percent of virtual memory used.
<b>BW Version</b>	The TIBCO BusinessWorks version currently in use on the server.
<b>Deployed Engines</b>	Total number of engines deployed on the server.
<b>Active Engines</b>	Number of engines currently active.
<b>Source</b>	Name of RTView Data Server sending this data (or localhost).
<b>Time Stamp</b>	Time this data was retrieved.

## All Servers Grid

This grid provides a list view of utilization metrics for all BW servers (represented in the All Servers Heatmap). Track and view in parallel the general performance of all BW servers. Click on a node to drill down to the ["Single Server Summary"](#) display and view detailed metrics for that particular server.



**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** , **Table** open commonly accessed displays.
- 6,047** The number of items currently in the display.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

Open the **Alert Views - RTView Alerts Table** display.

**Filter By:**

The display might include these filtering options:

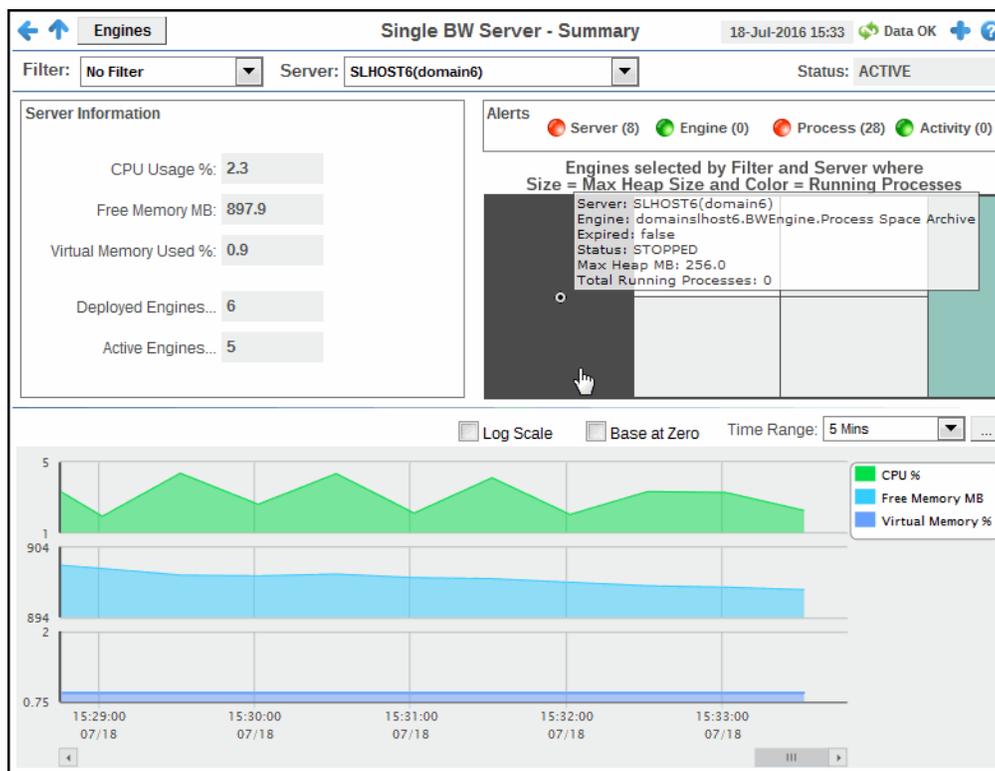
- Filter:** Choose a filter to show data for in the display. By default, the **Filter:** drop-down menu only contains the **No Filter** option. To create your own filtering options, see **Creating Customized Filters** in the User's Guide.
- Time Range** Choose a time range to show data for in the display. Options are: **All Data, 2 Mins, 5 Mins, 20 Mins, 1 Hour, 2 Hours, 4 Hours, 8 Hours, 24 Hours, 2 Days and 7 Days.**

**Fields and Data**

- Server Name** Name of the server.
- CPU Usage%** Percent of server CPU in use.
- Free Memory** Available physical memory (MB) remaining.
- Virtual Mem Used%** Percent of virtual memory used.
- State** Server status: ACTIVE or EXPIRED.
- Deployed Engines** Total number of engines deployed on the server.
- Active Engines** Number of engines currently active.
- Trend Graphs** Shows data for the server.
- CPU** Traces percent of server CPU in use.
  - MEM** Traces available physical memory remaining.
  - VMEM** Traces the percent of virtual memory used.

## Single Server Summary

Detailed performance metrics and alert status for a single BW server. Click on any alert indicator to drill down to the **BW Server - Tables** display to view current alerts for the selected server.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** **Table** open commonly accessed displays.
- 6,047** The number of items currently in the display.

- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

### Filter By:

The display might include these filtering options:

- Filter:** Choose a filter to show data for in the display. By default, the **Filter:** drop-down menu only contains the **No Filter** option. To create your own filtering options, see **Creating Customized Filters** in the User's Guide.
- Server:** Choose a server to see metrics for.
- Status** Server status: ACTIVE or EXPIRED.

### Server Information

- CPU Usage (%)** Percent of server CPU in use. Values are traced in trend graph (below).

<b>Free Memory (MB)</b>	Available physical memory remaining (in MB). Values are traced in trend graph (below).
<b>V. Memory Usage (%)</b>	Percent of virtual memory used. Values are traced in trend graph (below).
<b>Deployed Engines</b>	Number of engines currently active. Click to drill-down to details for deployed and active engines in the <a href="#">"All Engines Table" on page 1293</a> display.
<b>Active Engines</b>	Shows data for the server. Click to drill-down to details for active engines in the <a href="#">"All Engines Table" on page 1293</a> display.

### Alerts

Indicates the greatest severity level and the number of open **Server**, **Engine**, **Process**, and **Activity** alerts for the selected server. Values range from **0** to **2**, where **2** is the greatest Severity:

- One or more alerts exceeded their ALARM LEVEL threshold.
- One or more alerts exceeded their WARNING LEVEL threshold.
- No alert thresholds have been exceeded.

Click on the alert indicator to display the **BW Server - Tables** display, which contains a table listing the current alerts for the selected engine.

Time Stamp	Server	CPU Usage (%)	Free Memory (MB)	V. Memory Usage (%)	BW Version	Source	Deployed	Active	Expired	State
07/16/18 14:56:32	SLHOST16	1.29	1,520.20	18.47	v5.10	localhost	1	0		ACTIVE

Time	Alert Name	Alert Index	Alert Text

### Heatmap

Engines selected by Filter and Server, where the heatmap rectangle size = Max Heap Size and the heatmap rectangle color = Running Processes. Dark green is the highest value for the metric shown). Click on a node to drill down to a specific engine:

- Red indicates that the engine is expired.
- Gray indicates that the engine is stopped.

### Trend Graphs

Traces **CPU %**, **Free Memory MB** and **Virtual Memory %**.

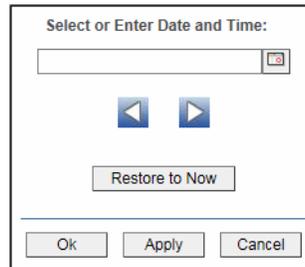
- Log Scale** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.

**Base at Zero**

Select to use zero (0) as the Y axis minimum for all graph traces.

**Time Range**

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Server Processes

Detailed information about operating system processes of a single BW Server. The heatmap shows server processes selected by Filter and Server, where the rectangle size equals memory usage and the rectangle color equals CPU percent usage.

NOTE: By default, this display is not enabled. For details, see **Enable BW Servers**.



**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- The number of items currently in the display.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

Open the **Alert Views - RTView Alerts Table** display.

**Filter By:**

The display might include these filtering options:

- Filter:** Choose a filter to show data for in the display. By default, the **Filter:** drop-down menu only contains the **No Filter** option. To create your own filtering options, see **Creating Customized Filters** in the User's Guide.
- Server:** Choose a server to see metrics for.

**Single Server Process - Summary**

Detailed information about a single operating system process running on a single BW Server.

NOTE: By default, this display is not enabled. For details, see **Enable BW Servers**.



**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** , **Table** open commonly accessed displays.
- 6,047** The number of items currently in the display.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

Open the **Alert Views - RTView Alerts Table** display.

**Filter By:**

The display might include these filtering options:

- Filter:** Choose a filter to show data for in the display. By default, the **Filter:** drop-down menu only contains the **No Filter** option. To create your own filtering options, see **Creating Customized Filters** in the User's Guide.
- Server:** Choose a server to see metrics for.
- Process:** Choose a server process.
- PID:** Choose a server PID.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .

The dialog box titled "Select or Enter Date and Time:" contains a text input field with a calendar icon on the right. Below the input field are two blue navigation arrows (left and right). A "Restore to Now" button is centered below the arrows. At the bottom of the dialog are three buttons: "Ok", "Apply", and "Cancel".

By default, the time range end point is the current time. To change the time range end point, click Calendar and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

---

## TIBCO BusinessWorks - HTML

The HTML version features an overview display, “[TIBCO BusinessWorks Overview Display - HTML](#)” (pictured below), and the following Views which can be found under **Components** tab > **Middleware**:

- “[BW Applications - HTML](#)”: The displays in this View present BusinessWorks application performance metrics.
- “[BW Containers - HTML](#)”: The displays in this View present BusinessWorks container performance metrics.
- “[BW Application Nodes - HTML](#)”: The displays in this View present BusinessWorks AppNode performance metrics.
- “[BW Application Slices - HTML](#)”: The displays in this View present BusinessWorks AppSlice performance metrics.
- “[BW Processes - HTML](#)”: The displays in this View present BusinessWorks process performance metrics.
- “[BW Activities - HTML](#)”: The displays in this View present BusinessWorks activity performance metrics.

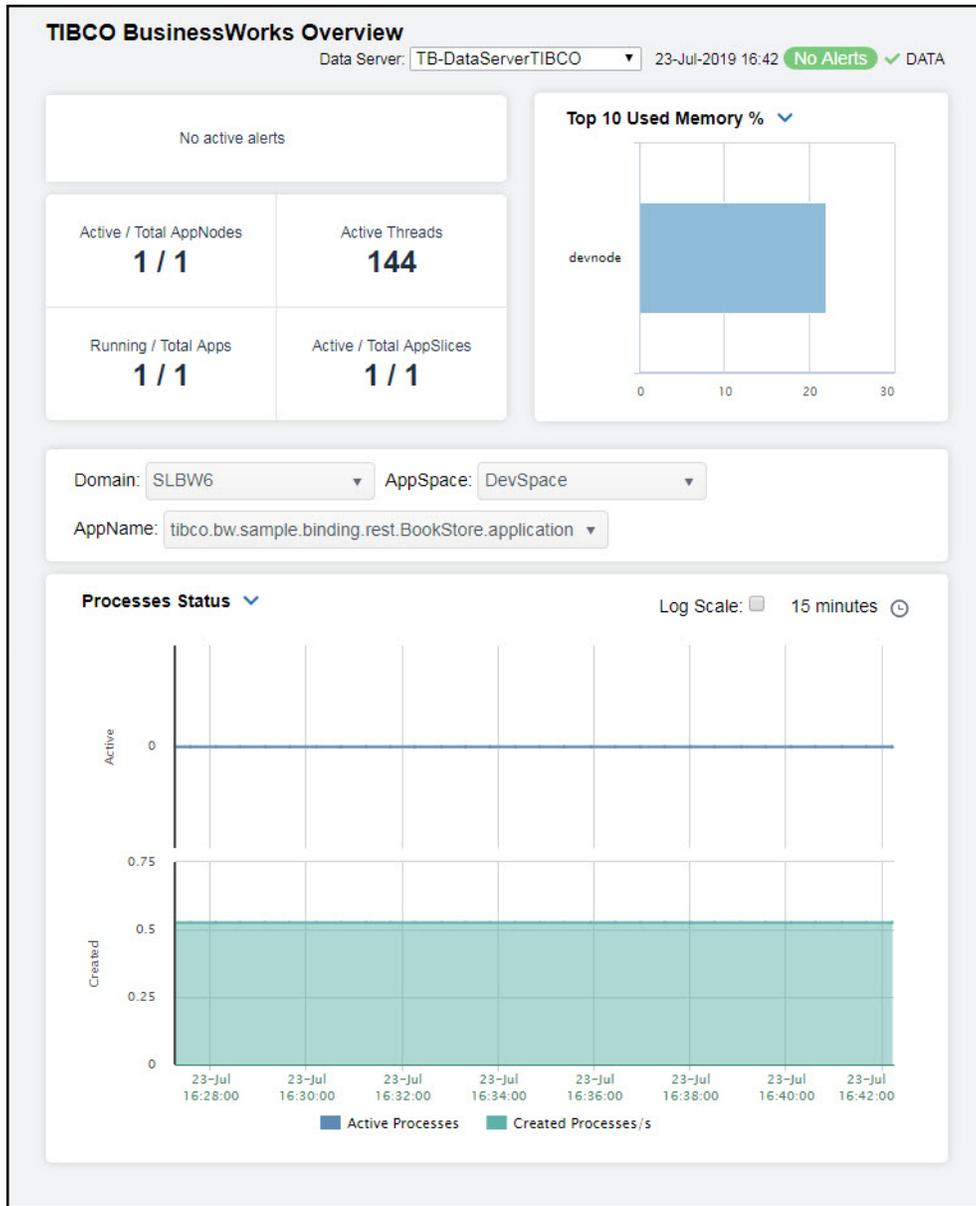
## TIBCO BusinessWorks Overview Display - HTML

The **TIBCO BusinessWorks Overview** is the top-level display for the TIBCO Enterprise BusinessWorks Monitor, which provides a good starting point for immediately getting the status of all your AppNodes, AppSlices, threads, and processes on your Data Server. You can select the RTView DataServer for which you want to see data and easily view the current data for that DataServer including:

- The total number of active alerts for the selected DataServer, including the total number of critical and warning alerts.
- The number of active AppNodes and the total number of AppNodes.
- The number of active threads on your connected DataServer.
- The number of running and total applications on your connected DataServer.
- The number of active and total AppSlices on your connected DataServer.
- A visual list of the top 10 servers containing the highest used CPUpercentage/used memory percentage/on your connected DataServer.

You can hover over each region in the upper half of the Overview to see more detail. You can also drill down to see even more detail by clicking on each respective region in the Overview. For example, clicking on the alerts in the **CRITICAL** and **WARNING** alerts region opens the **Alerts Table by Components** display.

The bottom half of the display provides a processes status and processes performance trend graph for a selected server. You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.



## BW Applications - HTML

These displays present process performance data for your BusinessWorks applications and AppSpaces across BusinessWorks Domains. Process metrics are totaled by application. Use these displays to monitor critical alerts for all your BusinessWorks applications, and investigate those alerts in lower-level displays. Clicking **BW Applications** from the left/navigation menu opens the ["TIBCO BusinessWorks Applications Table - HTML"](#) display, which shows all available utilization metrics for all BW applications. The options available under **BW Applications** are:

- **BW Applications Heatmap**: Opens the ["TIBCO BusinessWorks Applications Heatmap - HTML"](#), which shows server and alert status for all BW5 applications.
- **BW Application**: Opens the ["TIBCO BusinessWorks Application Summary - HTML"](#) display, which shows information for a single application.

## TIBCO BusinessWorks Applications Table - HTML

Investigate detailed utilization metrics for all BW applications. The **TIBCO BusinessWorks Applications Table** contains all metrics available for applications, including the number of active, failed, suspended, and created applications. Each row in the table contains data for a particular application. Choose a **Domain** and **AppSpace** from the drop-down menus to display activities for the selected Domain/AppSpace combination, or choose **All** from the drop downs to view all applications. Click a column header to sort column data in ascending or descending order. Double-click on a table row to drill-down to the ["TIBCO BusinessWorks Application Summary - HTML"](#) display and view metrics for that particular application. Toggle between the commonly accessed **Table** and **Heatmap** displays by clicking the drop down list on the display title.

**TIBCO BusinessWorks Applications Table** 18-Apr-2019 14:27 No Alerts DATA

Domain: - All - AppSpace: - All -

Applications: **3** Running: **3**

Domain	AppSpace	Application Name	Alert Level	Alert Count
SLBW6	DevSpace	SimpleTest.application	✓	0
SLBW6	DevSpace	tibco.bw.sample.binding.rest.BookStore.application	✓	0
SLBW6	Docker	tibco.bwce.sample.binding.rest.BookStore.applicatic	✓	0

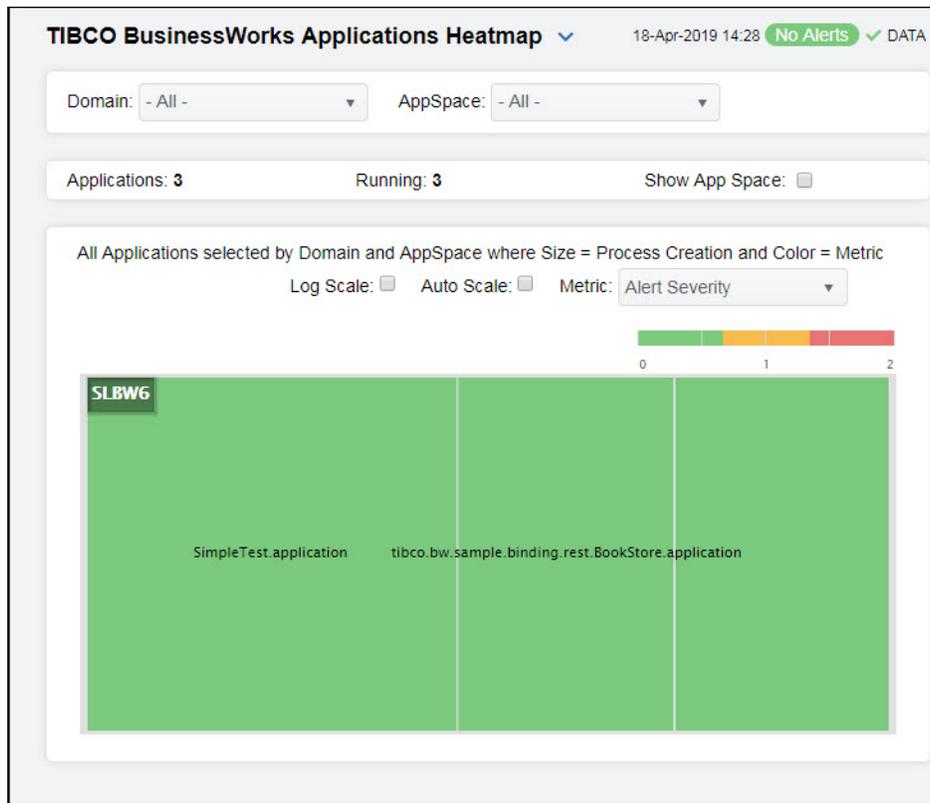
## TIBCO BusinessWorks Applications Heatmap - HTML

Clicking **BW Applications Heatmap** in the left/navigation menu opens the **TIBCO BusinessWorks Applications Heatmap**, which allows you to view the most critical BusinessWorks application alert states pertaining to process creation and execution for all nodes on which the applications are deployed. Use this display to quickly identify applications with critical alerts.

Each rectangle in the heatmap represents an application. The rectangle color indicates the most critical alert state associated with the application. The rectangle size represents process creation across applications; a larger size is a larger value.

Choose a domain and AppSpace from the drop-down menus. Choose a different metric to display from the **Metric** drop-down menu. Use the **Show App Space** check-box  to include or exclude labels in the heatmap. Mouse over a rectangle to see additional metrics. By default, this display shows **Alert Severity**.

Drill-down and investigate an application by clicking a rectangle in the heatmap to view details in the ["TIBCO BusinessWorks Application Summary - HTML"](#) display.



### Available Metrics

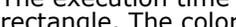
Select the metric driving the heatmap display. The default is **Alert Severity**. Each **Metric** has a color gradient bar that maps values to colors. The heatmap organizes the servers by domain, where each rectangle represents an application. Mouse-over any rectangle to display the current values of the metrics for the application. Click on a rectangle to drill-down to the associated ["TIBCO BusinessWorks Application Summary - HTML"](#) display for a detailed view of metrics for that particular application.

**Alert Severity** The maximum level of alerts in the heatmap rectangle. Values range from **0** - **2**, as indicated in the color gradient bar, where **2** is the highest Alert Severity:

- Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
- Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
- Green indicates that no metrics have exceeded their alert thresholds.

**Alert Count** The total number of critical and warning alerts in the heatmap rectangle. The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.

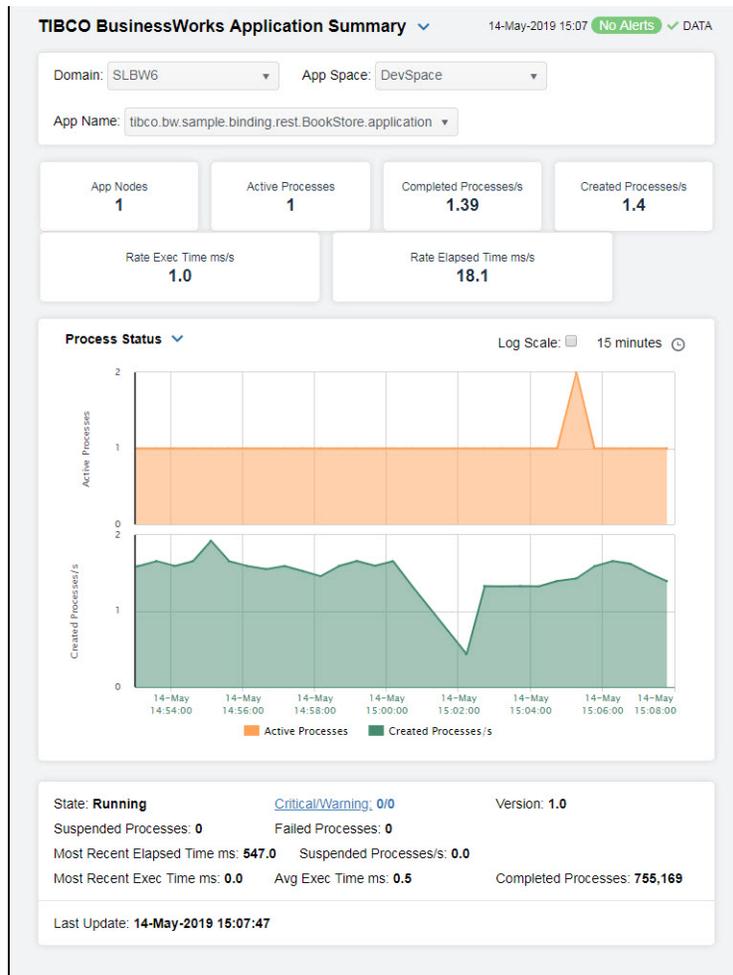
**Active Count** The total number of active processes in the heatmap rectangle. The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

<b>Completed Count</b>	The total number of completed processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Suspended Count</b>	The total number of suspended processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Failed Count</b>	The total number of failed processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Created / sec</b>	The number of processes created per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Suspended / sec</b>	The number of suspended processes per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Failed / sec</b>	The number of failed processes per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Exec Time / sec</b>	The process execution time per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Most Recent Exec Time</b>	The execution time for the most recently executed process in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Average Elapsed Time</b>	The average elapsed time for all processes in the heatmap rectangle, calculated by dividing the delta elapsed time for the interval by the delta completed, or the number of process instances that completed in the interval. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

## TIBCO BusinessWorks Application Summary - HTML

Clicking **BW Application** in the left/navigation menu opens the **TIBCO BusinessWorks Application Summary** display, which allows you to view current and historical metrics for a single BusinessWorks application across multiple nodes. Use this display to investigate performance issues of application AppNodes within an AppSpace. Use this display to view all available data for each AppNode by Domain and AppSpace.

Clicking on the information boxes at the top of the display takes you to the ["TIBCO BusinessWorks Application Nodes Table - HTML"](#) display or the ["TIBCO BusinessWorks Processes Table - HTML"](#) display, where you can view additional AppNode and Processes data. You can select from two different trend graphs: **Process Status** and **Process Performance**. In the **Process Status** trend graph region, you can view the created processes rate and number of active processes over a selected time range. In the Process Performance trend graph region, you can view the elapsed time rate and execution time rate over a selected time range. Clicking the **Critical/Warning** link at the bottom of the display opens the Alerts Table by Component display.



## BW Containers - HTML

These displays present process performance data for your BusinessWorks containers across BusinessWorks Domains. Use these displays to monitor critical alerts for all your BusinessWorks containers, and investigate those alerts in lower-level displays. Clicking **BW Containers** from the left/navigation menu opens the ["TIBCO BusinessWorks Containers Table - HTML"](#) display, which shows a tabular view of all available container performance data. The options available under **BW Containers** are:

- **BW Containers Heatmap**: Opens the ["TIBCO BusinessWorks Containers Heatmap - HTML"](#), which is a color-coded heatmap view of selected container performance metrics.
- **BW Container Summary**: Opens the ["TIBCO BusinessWorks Container Summary - HTML"](#) display, which shows current and historical metrics for a single container.

## TIBCO BusinessWorks Containers Table - HTML

This display provides a view of the most critical BusinessWorks container alert states pertaining to process creation and execution for all nodes on which the containers are deployed in a tabular format. Use this display to quickly identify containers with critical alerts. Each row in the table is a container in the selected domain.

By default, all containers are listed in the table, but you can enter a string in the **Container Name** filter field to limit the list of containers shown in the display. Click a column header to sort column data in numerical or alphabetical order. Click the **Running Only** check box to only view containers that are up and running in the table.

To view additional details for a specific container, drill-down and investigate by clicking the row in the table for the desired container, which opens the ["TIBCO BusinessWorks Container Summary - HTML"](#) display.

**TIBCO BusinessWorks Containers Table** 18-Apr-2019 14:31 No Alerts  DATA

Containers: 2      Running: 2      Running Only:

Container Name: \*       RegEx:

Application Name	Alert Level	Alert Count	State	Active	Contai
tibco.bwce.sample.binding.rest.BookStore.applicatio	✓	0	Running	0	✓
tibco.bwce.sample.binding.rest.BookStore.applicatio	✓	0	Running	0	✓

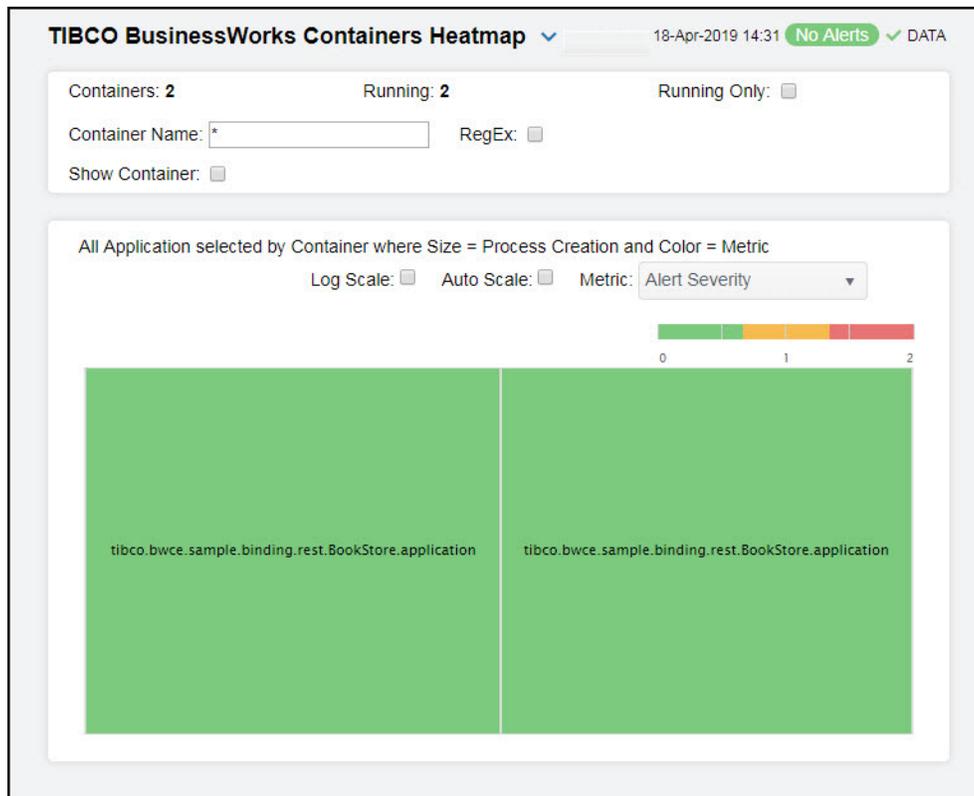
## TIBCO BusinessWorks Containers Heatmap - HTML

Clicking **BW Containers Heatmap** in the left/navigation menu opens the **TIBCO BusinessWorks Container Heatmap**, which allows you to view the most critical BusinessWorks container alert states pertaining to process creation and execution for all nodes on which the containers are deployed. Use this display to quickly identify containers with critical alerts.

Each rectangle in the heatmap represents a container. The rectangle color indicates the most critical alert state associated with the container. The rectangle size represents process creation across containers; a larger size is a larger value.

By default, all containers are listed in the heatmap, and is based on the **Alert Severity** metric. You can select a different metric from the **Metric** drop-down menu to view the heatmap based on a different metric. To view data shown for a specific container(s) in the display, enter a string in the **Container Name** filter field. Use the **Container Names** check-box  to include or exclude labels in the heatmap. You can mouse over a rectangle to see additional metrics. Toggle between the commonly accessed **Table** and **Heatmap** displays by clicking the drop down list on the display title.

Drill-down and investigate details for a specific container by clicking a rectangle in the heatmap, which opens the details for the selected container in the "[TIBCO BusinessWorks Container Summary - HTML](#)" display.



### Available Metrics

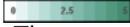
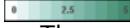
Select the metric driving the heatmap display. The default is **Alert Severity**. Each **Metric** has a color gradient bar that maps values to colors. The heatmap consists of multiple rectangles, where each rectangle represents a container. Mouse-over any rectangle to display the current values of the metrics for the container. Click on a rectangle to drill-down to the associated ["TIBCO BusinessWorks Container Summary - HTML"](#) display for a detailed view of metrics for that particular container.

**Alert Severity** The maximum level of alerts in the heatmap rectangle. Values range from **0** - **2**, as indicated in the color gradient  bar, where **2** is the highest Alert Severity:

-  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
-  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
-  Green indicates that no metrics have exceeded their alert thresholds.

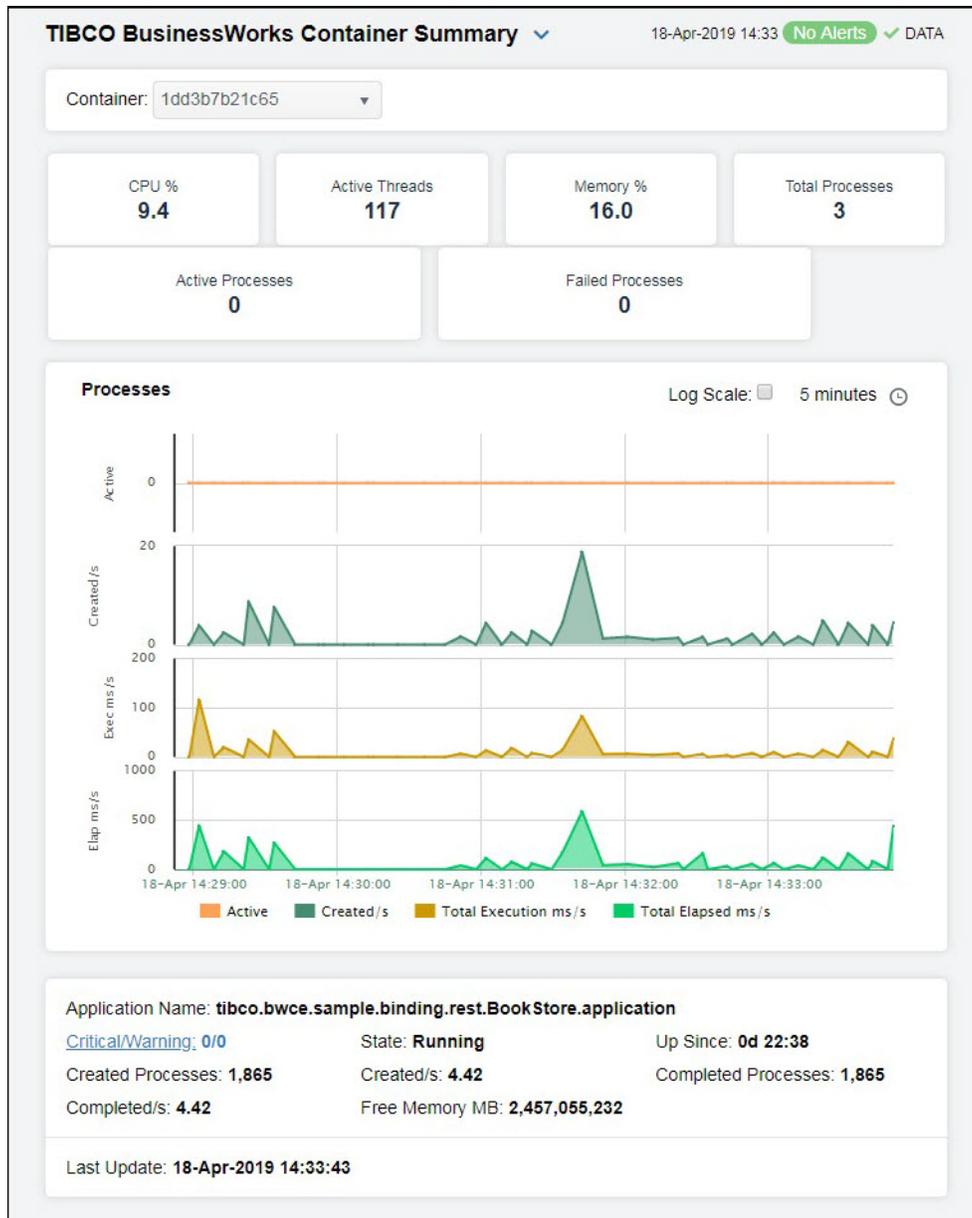
**Alert Count** The total number of critical and warning alerts in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.

**Active** The total number of active processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

<b>Completed</b>	The total number of completed processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Suspended</b>	The total number of suspended processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Failed</b>	The total number of failed processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Created/s</b>	The number of processes created per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Suspended/s</b>	The number of suspended processes per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Failed/s</b>	The number of failed processes per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Exec Time/s</b>	The process execution time per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Most Recent Exec Time</b>	The execution time for the most recently executed process in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Avg Exec Time</b>	The average execution time for all processes in the heatmap rectangle, calculated by dividing the delta execution time for the interval by the delta completed, or the number of process instances that completed in the interval. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
<b>Avg Elapsed Time</b>	The average elapsed time for all processes in the heatmap rectangle, calculated by dividing the delta elapsed time for the interval by the delta completed, or the number of process instances that completed in the interval. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

## TIBCO BusinessWorks Container Summary - HTML

Clicking **BW Container Summary** in the left/navigation menu opens the **TIBCO BusinessWorks Container Summary** display, which allows you to track utilization and performance metrics for specific BW containers. Clicking on the information boxes at the top of the display takes you to the "[TIBCO BusinessWorks Containers Table - HTML](#)" display, where you can view additional container data. In the **Processes** trend graph region, you can view the number of active containers, the container created rate, the total execution rate, and the total elapsed rate over a selected time range. Clicking the **Critical/Warning** link at the bottom of the display opens the Alerts Table by Component display.



## BW Application Nodes - HTML

These displays present internal JVM memory and host CPU utilization for BusinessWorks AppNodes and their resources, which can be useful because the AppNode performance is dependent on both internal and external factors and they sometimes interact. Clicking **BW Application Nodes** from the left/navigation menu opens the ["TIBCO BusinessWorks Application Nodes Table - HTML"](#) display, which shows a tabular view of all available utilization data. The options available under **BW Application Nodes** are:

- **BW Application Nodes Heatmap:** Opens the ["TIBCO BusinessWorks Application Nodes Heatmap - HTML"](#), which shows a color-coded heatmap view of utilization metrics.
- **BW Application Node:** Opens the ["TIBCO BusinessWorks Application Node Summary - HTML"](#) display, which shows Current and historical metrics for a single AppNode.

## TIBCO BusinessWorks Application Nodes Table - HTML

View BusinessWorks data shown in the ["TIBCO BusinessWorks Application Nodes Heatmap - HTML"](#) display, and additional details, in a tabular format. Use this display to view all available data for each AppNode by Domain and AppSpace. Each row in the table is an AppNode. Choose a domain and AppSpace from the drop-down menus. Click a column header to sort column data in numerical or alphabetical order. Drill-down and investigate by clicking a row to view details for the selected AppNode in the ["TIBCO BusinessWorks Application Node Summary - HTML"](#) display.

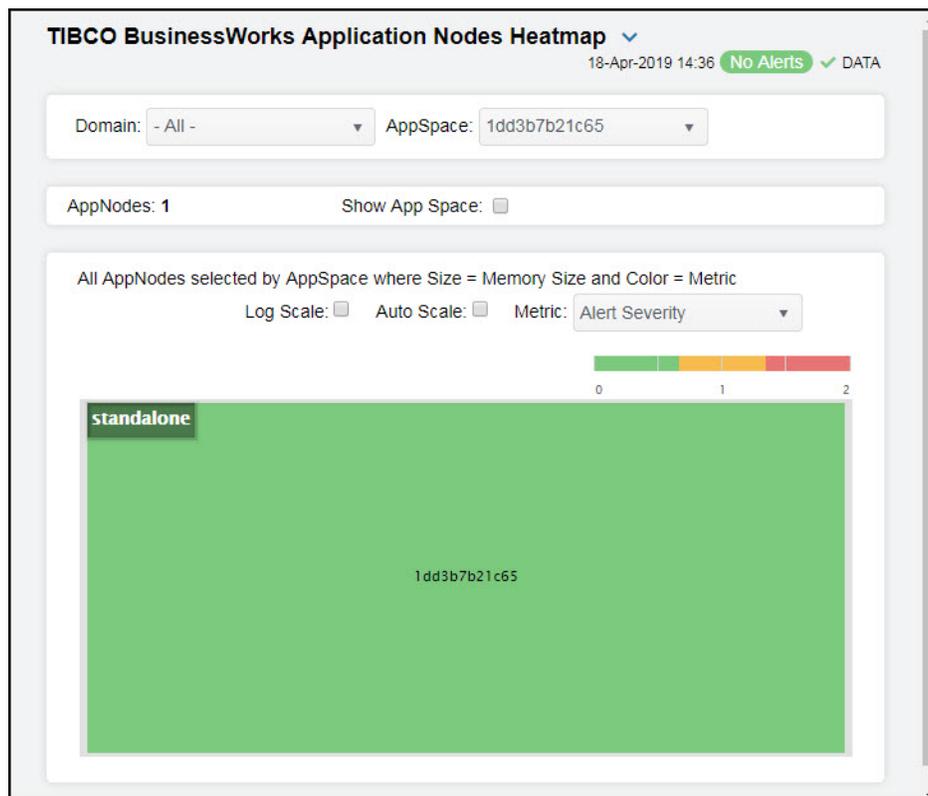
Domain	AppSpace	AppNode	Alert Level	Alert Count	State	Host
SLBW6	DevSpace	testnode	✓	0	ACTIVE	slhost9agent(SLBW6)
SLBW6	DevSpace	devnode	✓	0	ACTIVE	qawin5agent(SLBW6)

## TIBCO BusinessWorks Application Nodes Heatmap - HTML

Clicking **BW Application Nodes Heatmap** in the left/navigation menu opens the **TIBCO BusinessWorks Application Nodes Heatmap**, which allows you to view the most critical JVM memory and host resource utilization for BusinessWorks AppNodes. Use this display to quickly identify AppNodes with critical alerts.

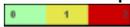
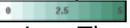
Each rectangle in the heatmap represents an AppNode. The rectangle color indicates the most critical alert state associated with the AppNode. The rectangle size represents the maximum memory used in the rectangle; a larger size is a larger value. Choose a domain and AppSpace from the drop-down menus. Choose a different metric to display from the **Metric** drop-down menu.

Use the **Show App Space** check-box  to include or exclude labels in the heatmap. Mouse over a rectangle to see additional metrics. By default, this display shows **Alert Severity**. Drill-down and investigate by clicking a rectangle in the heatmap to view details for the selected application in the ["TIBCO BusinessWorks Application Node Summary - HTML"](#) display. Toggle between the commonly accessed **Table** and **Heatmap** displays by clicking the drop down list on the display title.



### Available Metrics

Select the metric driving the heatmap display. The default is **Alert Severity**. Each **Metric** has a color gradient bar that maps values to colors. The heatmap organizes the servers by host, where each rectangle represents a server. Mouse-over any rectangle to display the current values of the metrics for the Server. Click on a rectangle to drill-down to the associated ["TIBCO BusinessWorks Application Node Summary - HTML"](#) display for a detailed view of metrics for that particular AppNode.

<b>Alert Severity</b>	<p>The maximum level of alerts in the heatmap rectangle. Values range from <b>0</b> - <b>2</b>, as indicated in the color gradient  bar, where <b>2</b> is the highest Alert Severity:</p> <ul style="list-style-type: none"> <li> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Alert Count</b>	<p>The total number of critical and warning alerts in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.</p>
<b>CPU Used%</b>	<p>The percent (%) CPU used in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum CPU used percentage in the heatmap. The middle value in the gradient bar indicates the average amount.</p>
<b>Memory Used%</b>	<p>The percent (%) memory used in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum memory used percentage in the heatmap. The middle value in the gradient bar indicates the average amount.</p>
<b>Active Processes</b>	<p>The number of currently active processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum number of active processes in the heatmap. The middle value in the gradient bar indicates the average count.</p>
<b>Created Processes</b>	<p>The number of processes created in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum number of created processes in the heatmap. The middle value in the gradient bar indicates the average count.</p>
<b>Created/sec</b>	<p>The number of processes created per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum creation rate in the heatmap. The middle value in the gradient bar indicates the average count.</p>
<b>Failed Processes</b>	<p>The number of failed processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum number of failed processes in the heatmap. The middle value in the gradient bar indicates the average number.</p>

## TIBCO BusinessWorks Application Node Summary - HTML

Clicking **BW Application Node** in the left/navigation menu opens the **TIBCO BusinessWorks Application Node Summary** display, which allows you to view current and historical utilization and performance metrics for a single BusinessWorks AppNode. Use this display to investigate performance issues on an AppNode. Clicking on the information boxes at the top of the display takes you to the ["TIBCO BusinessWorks Application Nodes Table - HTML"](#) display, where you can view additional AppNode data.

In the **Processes totals by AppNode and App** heatmap, you can select from the available metrics to view the current status of the processes running on the selected AppNode. Available metrics include **Alert Severity**, **Alert Count**, **Created/s**, and **Average Execution**.

You can select from two different trend graphs: **System Utilization** and **Memory Utilization**. In the **System Utilization** trend graph region, you can view the CPU percentage and number of threads over a selected time range. In the **Memory Utilization** trend graph region, you can view the memory percentage and number of bytes over a selected time range.

Clicking the **Critical/Warning** link at the bottom of the display opens the Alerts Table by Component display.



## BW Application Slices - HTML

These displays present process metrics totaled by Application and AppNode for AppSlices. This is useful to see how the application is distributed and how each part of it is performing. The AppSlice is the part of an application running on a specific AppNode when the application is deployed to multiple AppNodes. Clicking **BW Application Slices** from the left/navigation menu opens the "[TIBCO BusinessWorks Application Slices Table - HTML](#)" display, where each row in the table displays all available metrics for the AppSlice. The options available under **BW Application Slices** are:

- **BW Application Slices Heatmap**: Opens the "[TIBCO BusinessWorks Application Slices Heatmap - HTML](#)", which shows process execution metrics for all AppSlices.
- **BW Application Slice**: Opens the "[TIBCO BusinessWorks Application Slice Summary - HTML](#)" display, which shows current and historical metrics for a single AppSlice.

## TIBCO BusinessWorks Application Slices Table - HTML

Select a domain, AppSpace, and AppNode from the drop-down menus. Each row in the table is a different AppSpace and contains all metrics available for the AppSpace. You can limit the AppSlices listed in the table by entering a value in the **Application Name** filter field. By default, all AppSlices are listed in the table.

Click a column header to sort column data in ascending or descending order. Double-click on a table row to drill-down to the "[TIBCO BusinessWorks Application Slice Summary - HTML](#)" display and view metrics for that particular AppSlice. Toggle between the commonly accessed **Table** and **Heatmap** displays by clicking the drop down list on the display title.

**TIBCO BusinessWorks Application Slices Table** 18-Apr-2019 14:39 No Alerts DATA

Domain: - All - AppSpace: - All -  
 AppNode: - All -

Application Name:  AppSlices: **4** Running: **4**

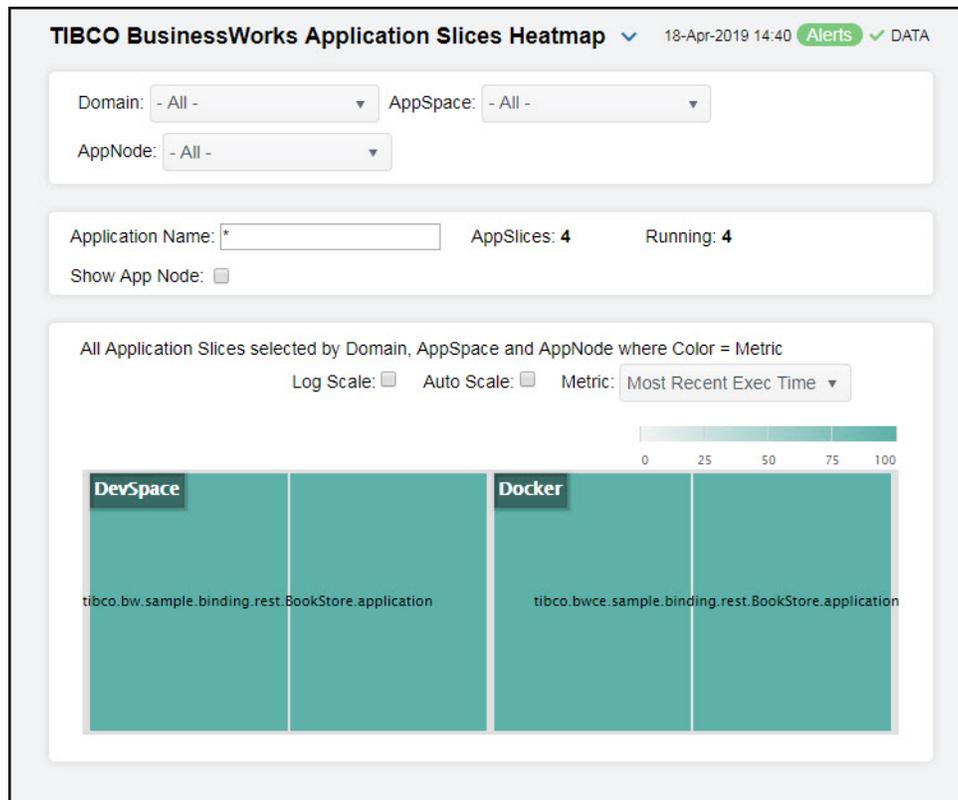
Domain	AppSpace	AppNode	Application Name
SLBW6	DevSpace	devnode	tibco.bw.sample.binding.rest.BookStore.application
SLBW6	DevSpace	testnode	SimpleTest.application
SLBW6	Docker	7c04e9140dde	tibco.bwce.sample.binding.rest.BookStore.application
SLBW6	Docker	7a5c3a46b942	tibco.bwce.sample.binding.rest.BookStore.application

## TIBCO BusinessWorks Application Slices Heatmap - HTML

Clicking **BW Application Slices Heatmap** in the left/navigation menu opens the **TIBCO BusinessWorks Application Slices Heatmap**, which allows you to view the most critical performance metrics for BusinessWorks AppSlices. Use this display to quickly identify AppSlices with high process execution numbers.

Each rectangle in the heatmap represents an AppSlice. The rectangle color indicates the process execution numbers for the AppSlice. The rectangle size represents the number of processes created in the rectangle; a larger size is a larger value. Move your mouse over a node to display current metrics.

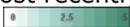
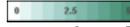
Choose a domain, AppSpace and AppNode from the drop-down menus. Enter a string in the **Application Name Filter** field to limit data shown in the display. Click the **Show AppNode** check-box  to include or exclude labels in the heatmap. Mouse over a rectangle to see additional metrics. By default, this display shows data based on the **Active Count** metric. Select a different metric from the Metric drop down menu to display the heatmap based on that metric. Drill-down and investigate by clicking a rectangle in the heatmap to view details for the selected application in the "[TIBCO BusinessWorks Application Slice Summary - HTML](#)" display.



### Available Metrics

Select the metric driving the heatmap display. The default is **Alert Severity**. Each **Metric** has a color gradient bar that maps values to colors. The heatmap is organized so that each rectangle represents an AppSlice. Mouse-over any rectangle to display the current values of the metrics for the AppSlice. Click on a rectangle to drill-down to the associated ["TIBCO BusinessWorks Application Slice Summary - HTML"](#) display for a detailed view of metrics for that particular AppSlice.

- Active Count** The total number of active processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
- Completed Count** The total number of completed processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum completed processes in an AppSlice in the heatmap. The middle value in the gradient bar indicates the average number of processes.
- Suspended Count** The total number of suspended processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum suspended processes in an AppSlice in the heatmap. The middle value in the gradient bar indicates the average number of suspended processes.
- Failed Count** The total number of failed processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum failed processes in an AppSlice in the heatmap. The middle value in the gradient bar indicates the average number of failed processes.

<b>Created/s</b>	The number of processes created per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum process creation rate in an AppSlice in the heatmap. The middle value in the gradient bar indicates the average creation rate.
<b>Suspended /s</b>	The number of suspended processes per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum process suspended rate in an AppSlice in the heatmap. The middle value in the gradient bar indicates the average suspended rate.
<b>Failed/s</b>	The number of failed processes per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum process failed rate in an AppSlice in the heatmap. The middle value in the gradient bar indicates the average failed rate.
<b>Exec Time/s</b>	The process execution time per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum execution time rate in an AppSlice in the heatmap. The middle value in the gradient bar indicates the average execution time rate.
<b>Most Recent Exec Time</b>	The execution time for the most recently executed process in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum execution time in an AppSlice in the heatmap. The middle value in the gradient bar indicates the average execution time.
<b>Average Exec Time</b>	The average execution time for all processes in the heatmap rectangle, calculated by dividing the delta execution time for the interval by the delta completed, or the number of process instances that completed in the interval. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum average execution time in the heatmap. The middle value in the gradient bar indicates the average execution time.
<b>Average Elapsed Time</b>	The average elapsed time for all processes in the heatmap rectangle, calculated by dividing the delta elapsed time for the interval by the delta completed, or the number of process instances that completed in the interval. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum average elapsed time in the heatmap. The middle value in the gradient bar indicates the average elapsed time.

## TIBCO BusinessWorks Application Slice Summary - HTML

Clicking **BW Application Slice** in the left/navigation menu opens the **TIBCO BusinessWorks Application Slice Summary** display, which allows you to view current and historical utilization and performance metrics for a single BusinessWorks AppSlice. Use this display to investigate performance issues on an AppSlice level. Choose a domain, AppSpace, AppNode, and Application Name from the drop-down menus.

Clicking on the information boxes at the top of the display takes you to the "[TIBCO BusinessWorks Application Slices Table - HTML](#)" display, where you can view additional AppSlice data.

In the **Process Data for AppNode organized by Application** heatmap, you can select from the available metrics to view the current status of the processes running on the selected AppNode. Available metrics include **Alert Severity**, **Alert Count**, **Created/s**, and **Average Execution**. Drill-down and investigate by clicking a AppSlice in the heatmap to view details in the "[TIBCO BusinessWorks Application Slice Summary - HTML](#)" display.

You can select from two different trend graphs: **Process Status** and **Process Performance**. In the **Process Status** trend graph region, you can view the number active processes and number of processes created per second over a selected time range. In the **Process Performance** trend graph region, you can view the process execution rate and the elapsed time rate over a selected time range.

Clicking the **Critical/Warning** link at the bottom of the display opens the Alerts Table by Component display.

### TIBCO BusinessWorks Application Slice Summary 14-May-2019 15:13 No Alerts DATA

Domain: SLBW6 AppSpace: DevSpace

AppNode: devnode

Application Name: tibco.bw.sample.binding.rest.BookStore.application

Processes  
**2**

Active Processes  
**1**

Completed Processes/s  
**1.92**

Created Processes/s  
**1.9**

Rate Exec Time ms/s  
**2.1**

Rate Elapsed Time ms/s  
**25.3**

Process Data for AppSlice organized by Application and AppNode where Color = Metric

Metric: Alert Severity

**Process Status** Log Scale:  15 minutes 🕒

State: <b>Running</b>	<a href="#">Critical/Warning: 0/0</a>	Version: <b>1.0</b>
Suspended Processes: <b>0</b>	Failed Processes: <b>0</b>	Average Elapsed Time ms: <b>14</b>
Suspended Processes/s: <b>0.0</b>	Failed Processes/s: <b>0.0</b>	Avg Exec Time ms: <b>0</b>
Completed: <b>755,727</b>		

Last Update: **14-May-2019 15:13:23**

## BW Processes - HTML

These displays present performance data for BusinessWorks processes. Use these displays to verify that individual BusinessWorks processes are executing and using resources as expected. Clicking **BW Processes** from the left/navigation menu opens the "[TIBCO BusinessWorks Processes Table - HTML](#)" display, where each row in the table displays all available metrics for the process. The options available under **BW Processes** are:

- **BW Processes Heatmap**: Opens the "[TIBCO BusinessWorks Processes Heatmap - HTML](#)", which shows process execution metrics for all processes.
- **BW Process**: Opens the "[TIBCO BusinessWorks Process Summary - HTML](#)" display, which shows current and historical metrics for a single process.

## TIBCO BusinessWorks Processes Table - HTML

Select a domain, AppSpace, AppNode, and Application from the drop-down menus. Each row in the table is a different process and contains all metrics available for the process.

Click a column header to sort column data in ascending or descending order. Double-click on a table row to drill-down to the "[TIBCO BusinessWorks Process Summary - HTML](#)" display and view metrics for that particular process. Toggle between the commonly accessed **Table** and **Heatmap** displays by clicking the drop down list on the display title.

**TIBCO BusinessWorks Processes Table** 18-Apr-2019 15:00 No Alerts DATA

Domain: SLBW6 AppSpace: Docker

AppNode: 7c04e9140dde

Application: tibco.bwce.sample.binding.rest.BookStore.application

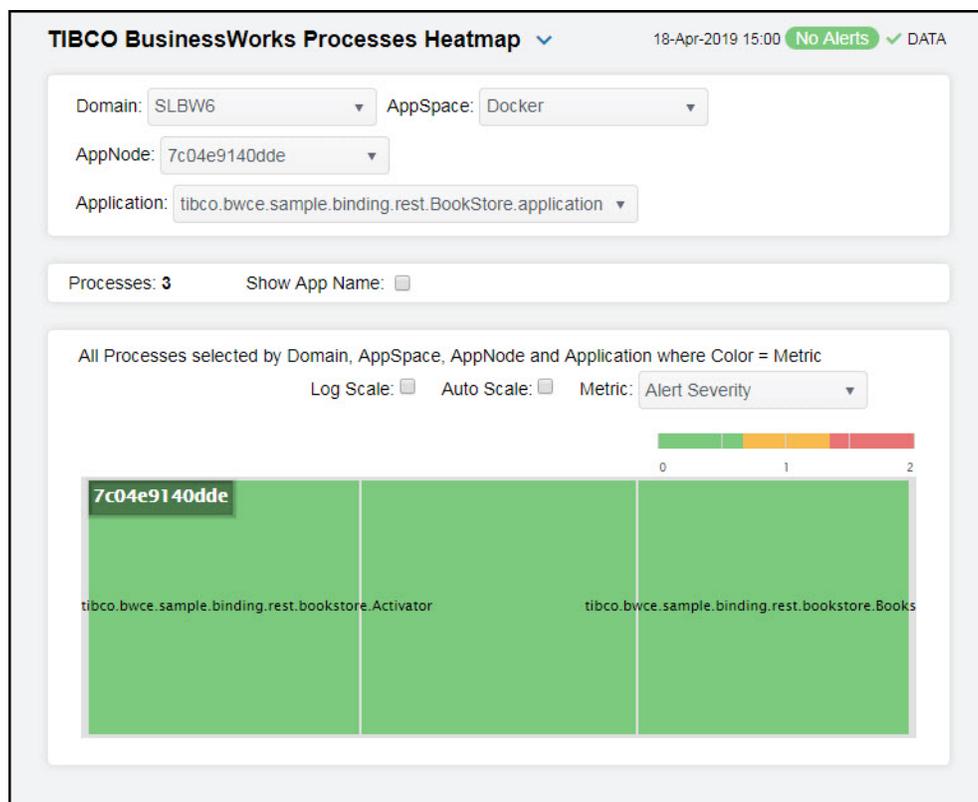
Processes: 3

Domain	AppSpace	AppNode	Application Name	P
SLBW6	Docker	7c04e9140dde	tibco.bwce.sample.binding.rest.BookStore.applica	tibco.bwce.sample
SLBW6	Docker	7c04e9140dde	tibco.bwce.sample.binding.rest.BookStore.applica	tibco.bwce.sample
SLBW6	Docker	7c04e9140dde	tibco.bwce.sample.binding.rest.BookStore.applica	tibco.bwce.sample

## TIBCO BusinessWorks Processes Heatmap - HTML

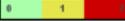
Clicking **BW Processes Heatmap** in the left/navigation menu opens the **TIBCO BusinessWorks Processes Heatmap**, which allows you to view the most critical BusinessWorks alerts pertaining to process creation and execution. Use this display to quickly identify processes with critical alerts.

The heatmap is organized by host with each rectangle representing a process. Move your mouse over a node to display current metrics. Choose a domain, AppSpace, AppNode, and Application from the drop-down menus. Click the **Show App Name** check-box  to include or exclude labels in the heatmap. Mouse over a rectangle to see additional metrics. By default, this display shows data based on the **Active Count** metric. Select a different metric from the **Metric** drop down menu to display the heatmap based on that metric. Drill-down and investigate by clicking a rectangle in the heatmap to view details for the selected application in the ["TIBCO BusinessWorks Process Summary - HTML"](#) display.



### Available Metrics

Select the metric driving the heatmap display. The default is **Alert Severity**. Each **Metric** has a color gradient bar that maps values to colors. The heatmap is organized so that each rectangle represents a process. Mouse-over any rectangle to display the current values of the metrics for the process. Click on a rectangle to drill-down to the associated ["TIBCO BusinessWorks Process Summary - HTML"](#) display for a detailed view of metrics for that particular process.

<b>Alert Severity</b>	<p>The maximum level of alerts in the heatmap rectangle. Values range from <b>0</b> - <b>2</b>, as indicated in the color gradient  bar, where <b>2</b> is the highest Alert Severity:</p> <ul style="list-style-type: none"> <li> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Alert Count</b>	<p>The total number of critical and warning alerts in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.</p>
<b>Active Count</b>	<p>The total number of active processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of active processes in the heatmap. The middle value in the gradient bar indicates the average count.</p>
<b>Completed Count</b>	<p>The total number of completed processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of completed processes in the heatmap. The middle value in the gradient bar indicates the average count.</p>
<b>Suspended Count</b>	<p>The total number of suspended processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of suspended processes in the heatmap. The middle value in the gradient bar indicates the average count.</p>
<b>Failed Count</b>	<p>The total number of failed processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of failed processes in the heatmap. The middle value in the gradient bar indicates the average count.</p>
<b>Created / sec</b>	<p>The number of processes created per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum number of processes created in the heatmap. The middle value in the gradient bar indicates the average number of created processes.</p>
<b>Suspended / sec</b>	<p>The number of suspended processes per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum number of suspended processes per second in the heatmap. The middle value in the gradient bar indicates the average rate of suspended processes.</p>
<b>Failed / sec</b>	<p>The number of failed processes per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum rate of failed processes in the heatmap. The middle value in the gradient bar indicates the average count.</p>
<b>Exec Time / sec</b>	<p>The process execution time per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum execution time rate in the heatmap. The middle value in the gradient bar indicates the average execution time rate.</p>
<b>Most Recent Exec Time</b>	<p>The execution time for the most recently executed process in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the most recent execution time in the heatmap. The middle value in the gradient bar indicates the average time.</p>

<b>Average Exec Time</b>	The average execution time for all processes in the heatmap rectangle, calculated by dividing the delta execution time for the interval by the delta completed, or the number of process instances that completed in the interval. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum average execution time in the heatmap. The middle value in the gradient bar indicates the average time.
<b>Most Recent Elapsed Time</b>	The elapsed time for the most recent process in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum most recent elapsed time in the heatmap. The middle value in the gradient bar indicates the average most recent elapsed time.
<b>Average Elapsed Time</b>	The average elapsed time for all processes in the heatmap rectangle, calculated by dividing the delta elapsed time for the interval by the delta completed, or the number of process instances that completed in the interval. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum average elapsed time in the heatmap. The middle value in the gradient bar indicates the average elapsed time.

## TIBCO BusinessWorks Process Summary - HTML

Clicking **BW Process** in the left/navigation menu opens the **TIBCO BusinessWorks Process Summary** display, which allows you to view current and historical execution metrics for a single BusinessWorks process. Choose a domain, AppSpace, AppNode, Application Name, and Process from the drop-down menus.

Clicking on the information boxes at the top of the display takes you to the "[TIBCO BusinessWorks Processes Table - HTML](#)" display, where you can view additional process data.

You can select from two different trend graphs: **Process Status** and **Process Performance**. In the **Process Status** trend graph region, you can view the number active processes and number of processes created per second over a selected time range. In the **Process Performance** trend graph region, you can view the process execution rate and the elapsed time rate over a selected time range.

Clicking the **Critical/Warning** link at the bottom of the display opens the Alerts Table by Component display.



## BW Activities - HTML

These displays present performance data for BusinessWorks activities. Use these displays to verify that individual BusinessWorks activities are executing and using resources as expected. Clicking **BW Activities** from the left/navigation menu opens the "[TIBCO BusinessWorks Activities Table - HTML](#)" display, where each row in the table displays all available metrics for the process. The options available under **BW Processes** are:

- **BW Activities Heatmap**: Opens the "[TIBCO BusinessWorks Activities Heatmap - HTML](#)", which shows process execution metrics for all activities.
- **BW Activity**: Opens the "[TIBCO BusinessWorks Activity Summary - HTML](#)" display, which shows current and historical metrics for a single activity.

## TIBCO BusinessWorks Activities Table - HTML

Select a domain, AppSpace, AppNode, AppName, and process from the drop-down menus. Each row in the table is a different activity and contains all metrics available for the activity.

Click a column header to sort column data in ascending or descending order. Double-click on a table row to drill-down to the "[TIBCO BusinessWorks Activity Summary - HTML](#)" display and view metrics for that particular activity. Toggle between the commonly accessed **Table** and **Heatmap** displays by clicking the drop down list on the display title.

**TIBCO BusinessWorks Activities Table** 16-May-2019 08:25 ✔ DATA

Domain: - All - AppSpace: - All -

AppNode: - All - AppName: - All -

Process: - All -

Activity Name:  Activities: 21

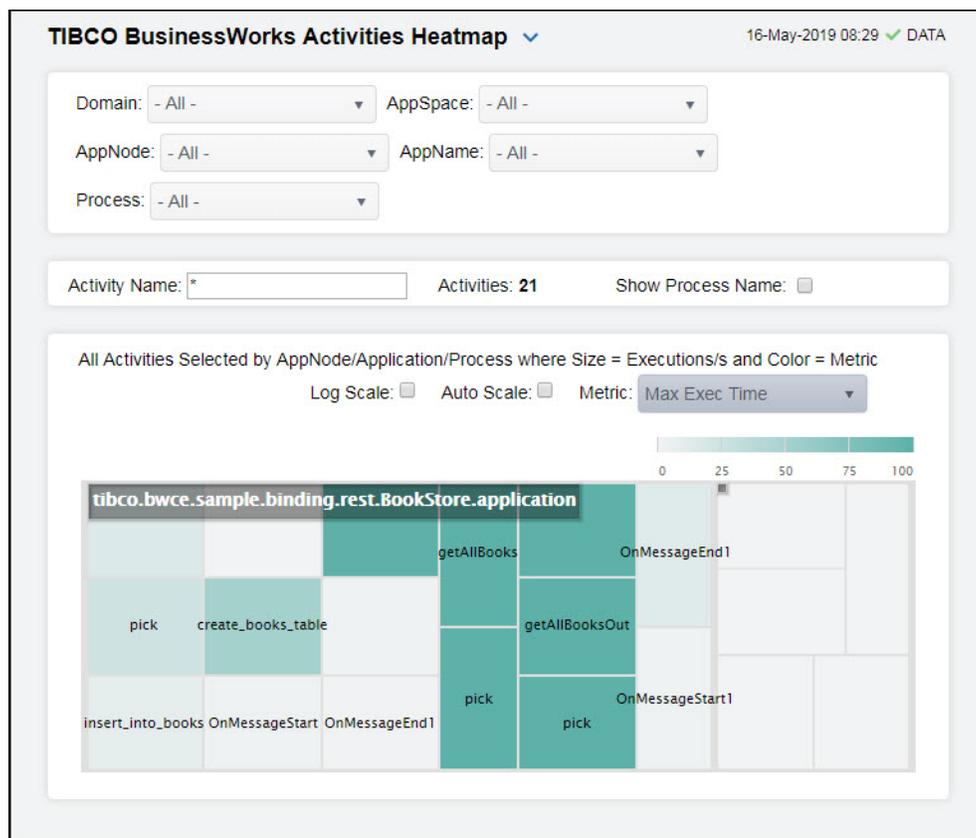
Domain	AppSpace	AppNode	Application ...	Process Name	Activity Name	Applicati
SLBW6	DevSpace	testnode				
standalone	2d89346a2f5c	2d89346a2f5c				
standalone	920523dc72fb	920523dc72fb				
SLBW6	DevSpace	centnode				
SLBW6	Docker	de968364dea9				
SLBW6	Docker	de968364dea9	tibco.bwce.samp	tibco.bwce.samp	getOut	1.0
SLBW6	Docker	de968364dea9	tibco.bwce.samp	tibco.bwce.samp	create_table_ev	1.0
SLBW6	Docker	de968364dea9	tibco.bwce.samp	tibco.bwce.samp	pick	1.0
SLBW6	Docker	de968364dea9	tibco.bwce.samp	tibco.bwce.samp	insert_into_book	1.0
SLBW6	Docker	de968364dea9	tibco.bwce.samp	tibco.bwce.samp	getBooks	1.0
SLBW6	Docker	de968364dea9	tibco.bwce.samp	tibco.bwce.samp	getAllBooksOut	1.0
SLBW6	Docker	de968364dea9	tibco.bwce.samp	tibco.bwce.samp	OnMessageEnd	1.0

## TIBCO BusinessWorks Activities Heatmap - HTML

Clicking **BW Activities Heatmap** in the left/navigation menu opens the **TIBCO BusinessWorks Activities Heatmap**, which allows you to view the most critical BusinessWorks alerts pertaining to activity creation and execution. Use this display to quickly identify activities with critical alerts.

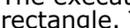
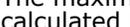
The heatmap is organized so that each rectangle represents a process. Move your mouse over a node to display current metrics. Click on a node to drill-down to the "[TIBCO BusinessWorks Activity Summary - HTML](#)" display to view specific metrics about process behavior over a specified period of time and determine which activity may be causing the bottleneck.

Choose a domain, AppSpace, AppNode, AppName, and process from the drop-down menus. Limit the activities listed in the table by specifying text in the **Activity Name** filter field. Click the **Show Process Name** check-box  to include or exclude labels in the heatmap. Mouse over a rectangle to see additional metrics. By default, this display shows data based on the **Error Count** metric. Select a different metric from the **Metric** drop down menu to display the heatmap based on that metric.



### Available Metrics

Select the metric driving the heatmap display. The default is **Error Count**. Each **Metric** has a color gradient bar that maps values to colors. The heatmap is organized so that each rectangle represents an activity. Mouse-over any rectangle to display the current values of the metrics for the activity. Click on a rectangle to drill-down to the associated "[TIBCO BusinessWorks Activity Summary - HTML](#)" display for a detailed view of metrics for that particular activity.

<b>Error Count</b>	The total number of errors in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of errors in the heatmap. The middle value in the gradient bar indicates the average error count.
<b>Errors/sec</b>	The number of errors per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of errors per second in the heatmap. The middle value in the gradient bar indicates the average errors per second count.
<b>Exec Count</b>	The total number of executions in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of executions in the heatmap. The middle value in the gradient bar indicates the average execution count.
<b>Exec Time/sec</b>	The process execution time per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average time.
<b>Most Recent Exec Time</b>	The execution time for the most recently executed activity in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average time.
<b>Max Exec Time</b>	The maximum execution time for all activities in the heatmap rectangle, calculated by dividing the delta execution time for the interval by the delta completed, or the number of process instances that completed in the interval. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average time.

## TIBCO BusinessWorks Activity Summary - HTML

Clicking **BW Activity** in the left/navigation menu opens the **TIBCO BusinessWorks Activity Summary** display, which allows you to view current and historical execution metrics for a single BusinessWorks activity. Choose a domain, AppSpace, AppNode, Application Name, Process, and Activity from the drop-down menus.

Clicking on the information boxes at the top of the display takes you to the "[TIBCO BusinessWorks Activities Table - HTML](#)" display, where you can view additional activity data.

You can select from two different trend graphs: **Utilization** and **Performance**. In the **Utilization** trend graph region, you can view the rate of executions and average number of errors over a selected time range. In the **Performance** trend graph region, you can view the average elapsed time and the average execution time over a selected time range.

### TIBCO BusinessWorks Activity Summary 16-May-2019 08:32 No Alerts DATA

Domain: SLBW6 AppSpace: Docker  
AppNode: de968364dea9  
AppName: tibco.bwce.sample.binding.rest.BookStore.application  
Process: tibco.bwce.sample.binding.rest.bookstore.Books  
Activity Name: getAllBooks

Executions/s <b>1.1</b>	Errors/s <b>0.0</b>	Rate Exec Time ms/s <b>0.5</b>	Avg Exec Time ms <b>0.5</b>
Rate Elapsed Time ms/s <b>38.8</b>		Avg Elapsed Time ms <b>36.3</b>	

#### Utilization Log Scale: 15 minutes

The utilization chart displays two metrics over a 15-minute period. The top chart shows 'Executions/s' as a green area chart, fluctuating between approximately 0.5 and 1.5. The bottom chart shows 'Avg Errors' as a red line, which remains constant at 0.0. The x-axis represents time from 08:18:00 to 08:32:00 on May 16, 2019.

Execution Count: <b>42,818</b>	Exec Count Since Reset: <b>0</b>	Current Executions: <b>10</b>
Error Count: <b>0</b>	Avg Errors: <b>0.0</b>	Current Errors: <b>0</b>
Min Execution Time: <b>0</b>	Current Exec Time: <b>5</b>	Max Execution Time: <b>2,011</b>
Min Elapsed Time: <b>0</b>	Current Elapsed Time: <b>363</b>	Max Elapsed Time: <b>17,317</b>

Last Update: 16-May-2019 08:32:04

---

## TIBCO BusinessWorks 5 Monitor - HTML

The following Views and their associated displays are in the Monitor. This section describes the Monitor displays and includes:

- ["TIBCO BusinessWorks 5 Overview Display"](#): Describes the TIBCO BusinessWorks Overview display.
- ["BW5 Servers - HTML"](#): The displays in this View present BusinessWorks 5.0 server performance metrics.
- ["BW5 Engines - HTML"](#): The displays in this View present BusinessWorks 5.0 engine performance metrics.
- ["BW5 Processes - HTML"](#): The displays in this View present BusinessWorks 5.0 process performance metrics.
- ["BW5 Activities - HTML"](#): The displays in this View present BusinessWorks 5.0 activity performance metrics.

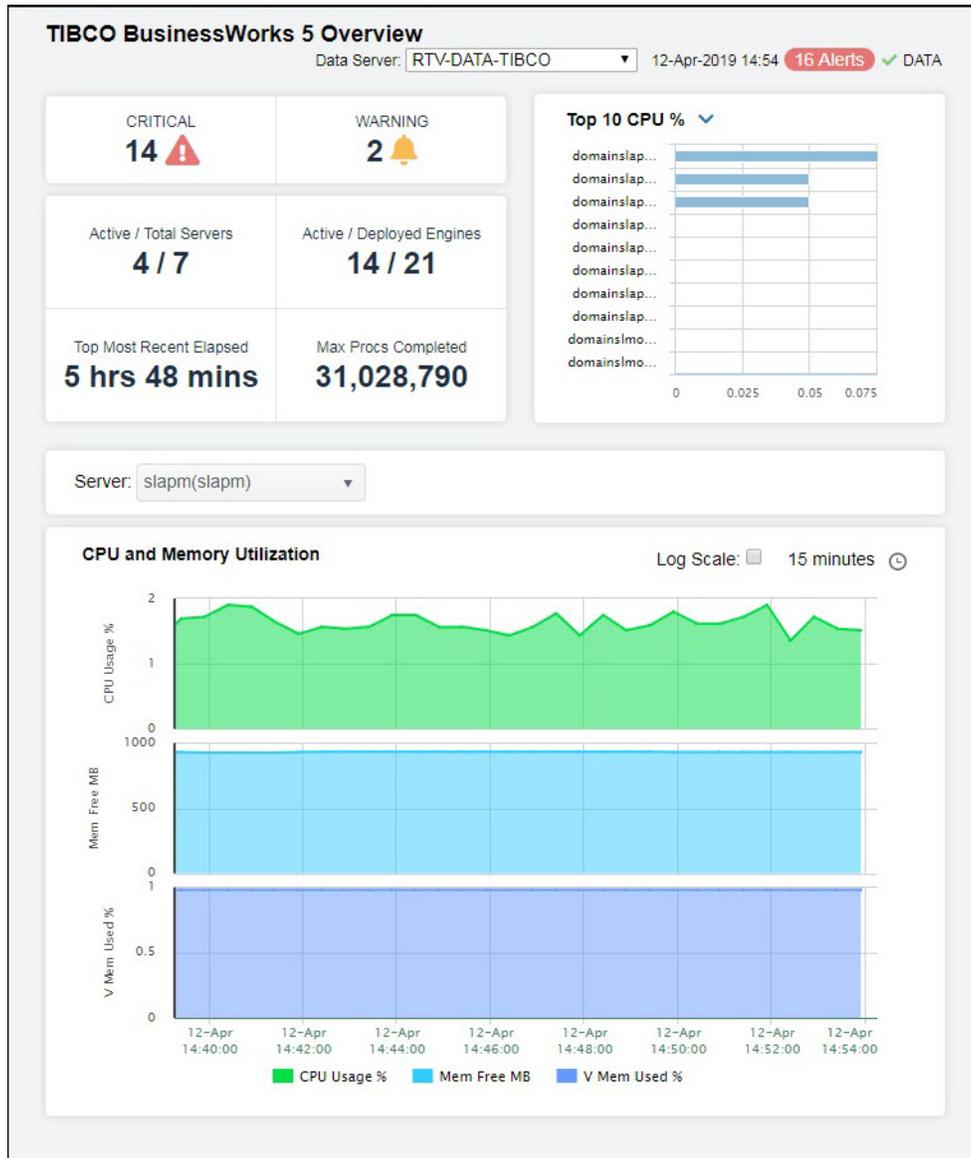
### TIBCO BusinessWorks 5 Overview Display

The **TIBCO BusinessWorks 5 Overview** is the top-level display for the TIBCO Enterprise BusinessWorks 5 Monitor, which provides a good starting point for immediately getting the status of all your servers, engines, and processes on your Data Server. You can select the RTView DataServer for which you want to see data and easily view the current data for that DataServer including:

- The total number of active alerts for the selected DataServer, including the total number of critical and warning alerts.
- The number of active servers and the total number of servers.
- The number of active and deployed engines on your connected DataServer.
- The top most recent elapsed time for a process on your connected DataServer.
- The maximum number of processes completed on one engine on your connected DataServer.
- A visual list of the top 10 servers containing the highest CPU usage percentage/memory used percentage/completed processes/error rate per second on your connected DataServer.

You can hover over each region in the upper half of the Overview to see more detail. You can also drill down to see even more detail by clicking on each respective region in the Overview. For example, clicking on the alerts in the **CRITICAL** and **WARNING** alerts region opens the **Alerts Table by Components** display.

The bottom half of the display provides a CPU and memory utilization trend graph for a selected server. You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.



## BW5 Servers - HTML

These displays present performance metrics and alert status for all BW5 servers. Clicking **BW5 Servers** from the left/navigation menu opens the “[TIBCO BusinessWorks 5 Servers Table - HTML](#)” display, which shows all available utilization metrics for all BW5 servers. The options available under **BW5 Servers** are:

- **All Servers Heatmap**: Opens the “[TIBCO BusinessWorks 5 Servers Heatmap - HTML](#)”, which shows server and alert status for all BW5 servers.
- **BW5 Server Summary**: Opens the “[TIBCO BusinessWorks 5 Server Summary - HTML](#)” display, which shows information for a single BW5 server.

## TIBCO BusinessWorks 5 Servers Table - HTML

Investigate detailed utilization metrics for all BW servers. The **TIBCO BusinessWorks 5 Servers Table** contains all metrics available for servers, including CPU usage percentage, free memory, and percentage of virtual memory used. Each row in the table contains data for a particular server. Click a column header to sort column data in ascending or descending order. Double-click on a table row to drill-down to the “[TIBCO BusinessWorks 5 Server Summary - HTML](#)” display and view metrics for that particular server. Toggle between the commonly accessed **Table** and **Heatmap** displays by clicking the drop down list on the display title.

**TIBCO BusinessWorks 5 Servers Table** ▾ 15-Apr-2019 13:37 16 Alerts ✓ DATA

Servers: 7

Server	Expired	Alert Level	Alert Count	CPU Usage %	Mem Free MB	VM Use
SLHOST21(dev)	🕒	⚠️	9	15.3	3,311.8	
SLHOST16(sl_qa_conn)	🕒	⚠️	2	74.7	1,332.0	
SLHOST17(sl_amx)	🕒	⚠️	2	49.9	323.1	
slxp10(slapm)		✅	0	1.5	1,162.6	
slapm(slapm)		⚠️	2	2.0	918.6	
slhpux11(simon)		✅	0	0.0	975.7	
slsl4-64(simon)		⚠️	1	67.5	101.0	

## TIBCO BusinessWorks 5 Servers Heatmap - HTML

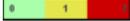
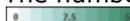
Clicking **All Servers Heatmap** in the left/navigation menu opens the **TIBCO BW5 Servers Heatmap**, which allows you to view the status and alerts of all BW5 servers. Use the **Metric** drop-down menu to view the **Alert Severity**, **Alert Count**, **CPU Used Percentage**, **Virtual Memory Used Percentage**, **Free Memory**, **Deployed Engines**, or **Active Engines**.

The heatmap is organized by host with each rectangle representing a server. The rectangle color indicates the most critical alert state. Click on a node to drill-down to the ["TIBCO BusinessWorks 5 Server Summary - HTML"](#) display and view metrics for a particular server. Toggle between the commonly accessed **Table** and **Heatmap** displays by clicking the drop down list on the display title. Mouse-over rectangles to view more details about host performance and status.



### Available Metrics

Select the metric driving the heatmap display. The default is **Alert Severity**. Each **Metric** has a color gradient bar that maps values to colors. The heatmap organizes the servers by host, where each rectangle represents a server. Mouse-over any rectangle to display the current values of the metrics for the Server. Click on a rectangle to drill-down to the associated ["TIBCO BusinessWorks 5 Server Summary - HTML"](#) display for a detailed view of metrics for that particular server.

<b>Alert Severity</b>	<p>The maximum level of alerts in the heatmap rectangle. Values range from <b>0</b> - <b>2</b>, as indicated in the color gradient  bar, where <b>2</b> is the highest Alert Severity:</p> <ul style="list-style-type: none"> <li> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Alert Count</b>	<p>The total number of critical and warning alerts in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.</p>
<b>CPU Used%</b>	<p>The percent (%) CPU used in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the alert threshold of <b>BwServersCpuUsedHigh</b>, which is <b>100</b>. The middle value in the gradient bar indicates the middle value of the range (the default is <b>50</b>).</p>
<b>V(irtual) Memory Used%</b>	<p>The percent (%) virtual memory used in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the alert threshold of virtual memory used percentage in the heatmap. The middle value in the gradient bar indicates the middle value of the range (the default is <b>50</b>).</p>
<b>Free Memory</b>	<p>The amount of free memory in the heatmap rectangle, in megabytes. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the amount of free memory in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>Deployed Engines</b>	<p>The number of deployed engines in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the number of deployed engines in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>Active Engines</b>	<p>The number of active engines in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the number of active engines in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>

## TIBCO BusinessWorks 5 Server Summary - HTML

Clicking **BW5 Server Summary** in the left/navigation menu opens the **TIBCO BusinessWorks 5 Server Summary** display, which allows you to track utilization and performance metrics for specific BW5 servers. Clicking on the information boxes at the top of the display takes you to the ["TIBCO BusinessWorks 5 Servers Table - HTML"](#) display, where you can view additional servers data. In the **CPU and Memory Utilization** trend graph region, you can view CPU usage percentage, free memory, and virtual memory used percentage over a selected time range. Clicking the **Critical/Warning** link at the bottom of the display opens the Alerts Table by Component display.



## BW5 Engines - HTML

These displays present performance metrics and alert status for all BW5 engines. Clicking **BW5 Engines** from the left/navigation menu opens the ["TIBCO BusinessWorks 5 Engines Table - HTML"](#) display, which shows all available utilization metrics for all BW5 engines. The options available under **BW5 Engines** are:

- **All Engines Heatmap**: Opens the ["TIBCO BusinessWorks 5 Engines Heatmap - HTML"](#), which shows engine and alert status for all BW5 servers.
- **BW5 Engine Summary**: Opens the ["TIBCO BusinessWorks 5 Engine Summary - HTML"](#) display, which shows information for a single BW5 engine.

## TIBCO BusinessWorks 5 Engines Table - HTML

Investigate detailed utilization metrics for all BW engines. The **TIBCO BusinessWorks 5 Engines Table** contains all metrics available for engines, including memory usage, memory used percentage, and CPU used percentage. You can enter a string in the **Filter by Engine Name** field to show only engines in the table with names that contain the string. For example, if you enter the string Madrid, all engines with Madrid in the engine name are shown in the table. If no entry is made, all engine names are shown. For most use cases, you can enter a portion of the engine name. Each row in the table contains data for a particular engine. Click a column header to sort column data in ascending or descending order. Double-click on a table row to drill-down to the ["TIBCO BusinessWorks 5 Engine Summary - HTML"](#) display and view metrics for that particular engine. Toggle between the commonly accessed **Table** and **Heatmap** displays by clicking the drop down list on the display title.

**Note:** Metrics are made available by the Hawk microagent for the engine (for details, refer to documentation for TIBCO BusinessWorks Administration, Appendix A: TIBCO Hawk Microagent Methods).

**TIBCO BusinessWorks 5 Engines Table** 15-Apr-2019 14:20 2 Alerts DATA

Server:

Filter by Engine Name:  Active Only:

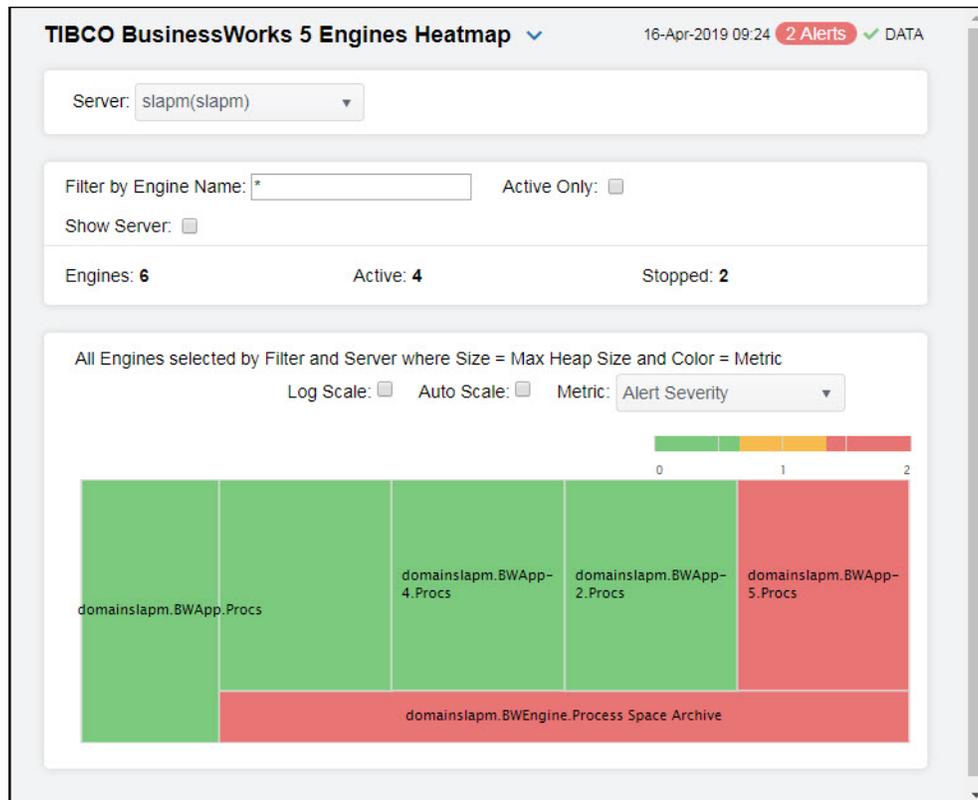
Engines: **6** Active: **4** Stopped: **2**

Server	Engine Name	Alert Level	Alert Count
slapm(slapm)	domainslapm.BWApp-3.Procs	✓	0
slapm(slapm)	domainslapm.BWApp-4.Procs	✓	0
slapm(slapm)	domainslapm.BWApp.Procs	✓	0
slapm(slapm)	domainslapm.BWApp-2.Procs	✓	0
slapm(slapm)	domainslapm.BWApp-5.Procs	⚠	1
slapm(slapm)	domainslapm.BWEngine.Process Space Arch	⚠	1

## TIBCO BusinessWorks 5 Engines Heatmap - HTML

Clicking **All Engines Heatmap** in the left/navigation menu opens the **TIBCO BW5 Engines Heatmap**, which allows you to view the status and alerts of all BW5 engines. You can enter a string in the **Filter by Engine Name** field to show only engines in the heatmap with names that contain the string. For example, if you enter the string Madrid, all engines with Madrid in the engine name are shown in the table. If no entry is made, all engine names are shown. For most use cases, you can enter a portion of the engine name. Use the **Metric** drop-down menu to view the **Alert Severity**, **Alert Count**, **CPU Used Percentage**, **Memory Used Percentage**, **Running Processes**, or **Error Count**.

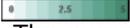
The heatmap is organized by host with each rectangle representing an engine. Rectangle size represents Max Heap Size and the color indicates the most critical alert state. Click on a node to drill-down to the ["TIBCO BusinessWorks 5 Engine Summary - HTML"](#) display and view metrics for a particular engine. Toggle between the commonly accessed **Table** and **Heatmap** displays by clicking the drop down list on the display title. Mouse-over rectangles to view more details about host performance and status.



### Available Metrics

Select the metric driving the heatmap display. The default is **Alert Severity**. Each **Metric** has a color gradient bar that maps values to colors. The heatmap organizes the servers by host, where each rectangle represents an engine. Mouse-over any rectangle to display the current values of the metrics for the engine. Click on a rectangle to drill-down to the associated ["TIBCO BusinessWorks 5 Engine Summary - HTML"](#) display for a detailed view of metrics for that particular engine.

- Alert Severity** The maximum level of alerts in the heatmap rectangle. Values range from **0** - **2**, as indicated in the color gradient bar, where **2** is the highest Alert Severity:
- Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
  - Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
  - Green indicates that no metrics have exceeded their alert thresholds.
- Alert Count** The total number of critical and warning alerts in the heatmap rectangle. The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.
- CPU Used%** The percent (%) CPU used in the heatmap rectangle. The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the alert threshold of **BwEngineCpuUsedHigh**, which is **100**. The middle value in the gradient bar indicates the middle value of the range (the default is **50**).

<b>Memory Used%</b>	The percent (%) memory used in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the alert threshold of <b>BwEngineMemUsedHigh</b> , which is <b>100</b> . The middle value in the gradient bar indicates the middle value of the range (the default is <b>50</b> ).
<b>Active Processes</b>	The number of currently active processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of active processes in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
<b>Running Processes</b>	The number of currently running processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of running processes in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
<b>Created Processes</b>	The number of created processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of created processes in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
<b>Created/sec</b>	The number of created processes in the heatmap rectangle, per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum rate of processes created per second in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
<b>Error Count</b>	The total number of errors in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of errors in the heatmap. The middle value in the gradient bar indicates the average alert count.

## TIBCO BusinessWorks 5 Engine Summary - HTML

Clicking **BW5 Engine Summary** in the left/navigation menu opens the **TIBCO BusinessWorks 5 Engine Summary** display, which allows you to track utilization and performance metrics for specific BW5 engines and their associated processes. Clicking on the information boxes at the top of the display takes you to the "[TIBCO BusinessWorks 5 Engines Table - HTML](#)" display or to the "[TIBCO BusinessWorks 5 Processes Table - HTML](#)" display (depending on which box you select), where you can view additional data on engines and processes. In the **Processes organized by Server/Engine** heatmap, you can select from the available metrics to view the current status of the processes running on the selected engine. Available metrics include **Alert Severity, Alert Count, Created Processes, Completed Processes, and Average Execution**.

There are two options in the trend graph: **System Utilization** and **Memory Utilization**. In the **System Utilization** option on the trend graph, you can view trend data for running processes and CPU percentage over a selected time range. In the **Memory Utilization** option on the trend graph, you can view trend data for used memory (in megabytes) and current memory (in megabytes) over a selected time range.

Clicking the **Critical/Warning** link at the bottom of the display opens the Alerts Table by Component display.

### TIBCO BusinessWorks 5 Engine Summary 16-Apr-2019 10:23 No Alerts DATA

Server: sle14-64(simon) Engine: domainsimon.BWApp-10.Procs

CPU %  
**0.0**

Total Threads  
**8**

Memory Used %  
**20.0**

Total Processes  
**11**

Running Processes  
**22**

Total Errors  
**0**

Processes organized by Server/Engine where Color = Metric

Metric: Alert Severity

0
  1
  2

process02.process	process06.process	process07.process	process04.process
process01.process	process05.process	process08.process	process09.process

**Memory Utilization** Log Scale:  5 minutes ⌵

The top chart shows Current Mem MB (light blue area) which is constant at approximately 226.74 MB. The bottom chart shows Used Mem MB (dark blue area) which starts at ~130 MB, drops to ~40 MB at 10:31:00, and then slightly increases to ~44.45 MB at 10:32:15.

State: **ACTIVE** [Critical/Warning: 0/0](#) Uptime: 1 yr 209 days

Free Memory MB: **203.58** Domain: **domainsimon** BW Version:

Total Created Processes: **17,626,882** Total Completed Processes: **17,567,958**

Total Aborted Processes: **58,902**

---

Last Update: **16-Apr-2019 10:34:16**

## BW5 Processes - HTML

These displays present performance metrics and alert status for all BW5 processes. Clicking **BW5 Processes** from the left/navigation menu opens the ["TIBCO BusinessWorks 5 Processes Table - HTML"](#) display, where each row in the table displays all available metrics from the Hawk microagent for a process. The options available under **BW5 Processes** are:

- **All Processes Heatmap**: Opens the ["TIBCO BusinessWorks 5 Processes Heatmap - HTML"](#), which shows process execution metrics for all BW Engines.
- **BW5 Process Summary**: Opens the ["TIBCO BusinessWorks 5 Process Summary - HTML"](#) display, which shows historical and current metrics for a single process, including average execution times and execution counts.

## TIBCO BusinessWorks 5 Processes Table - HTML

Select a server and engine from the drop-down menus. Each row in the table is a different process. The table displays all metrics available from the Hawk microagent for a process. (Refer to documentation for TIBCO BusinessWorks Administration, see Appendix A: TIBCO Hawk Microagent Methods).

Click a column header to sort column data in ascending or descending order. Double-click on a table row to drill-down to the ["TIBCO BusinessWorks 5 Process Summary - HTML"](#) display and view metrics for that particular process. Toggle between the commonly accessed **Table** and **Heatmap** displays by clicking the drop down list on the display title.

**TIBCO BusinessWorks 5 Processes Table** ▼ 16-Apr-2019 11:00 4 Alerts ✓ DATA

Server: sle14-64(simon) Engine: domainsimon.BWApp-10.Procs

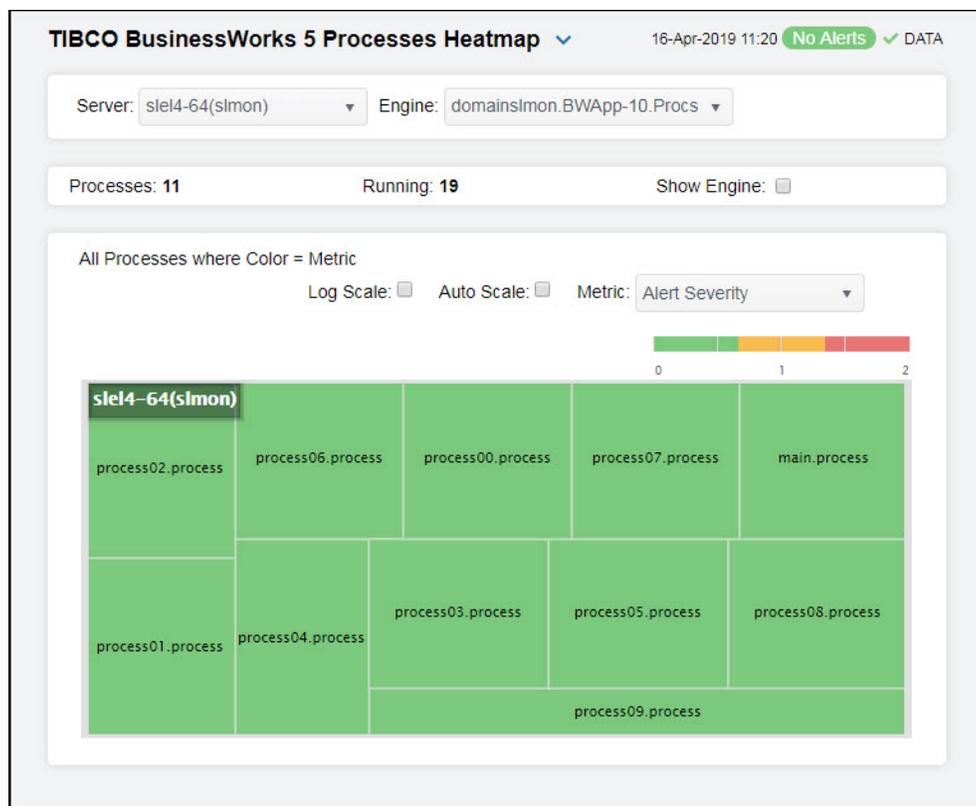
Processes: 11 Running: 27

Process	Alert Level	Alert Count	CPU %	Created/s	Cor
process02.process	▲	1	0.0	0.03	
process01.process	▲	1	0.0	0.03	
process06.process	✓	0	0.0	0.03	
process00.process	✓	0	0.0	0.03	
process07.process	✓	0	0.0	0.03	
main.process	✓	0	0.0	0.03	
process04.process	▲	1	0.0	0.03	
process03.process	▲	1	0.0	0.03	
process05.process	✓	0	0.0	0.03	
process08.process	✓	0	0.0	0.03	
process09.process	✓	0	0.0	0.03	

## TIBCO BusinessWorks 5 Processes Heatmap - HTML

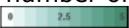
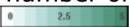
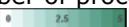
Clicking **All Processes Heatmap** in the left/navigation menu opens the **TIBCO BW5 Processes Heatmap**, which allows you to view the status and alerts of all BW5 processes for all engines or for a specific engine. Use the **Metric** drop-down menu to view processes in the heatmap by the **Alert Severity, Alert Count, CPU Used Percentage, Completed Processes, Active Processes, Aborted Processes, Suspended Processes, Execution Time per second, Created per second, Aborted per second, Suspended per second, Most Recent Execution Time, Average Execution Time, Most Recent Elapsed Time, and Average Elapsed Time.**

The heatmap is organized by host with each rectangle representing a process. Move your mouse over a node to display current metrics. Click on a node to drill-down to the ["TIBCO BusinessWorks 5 Process Summary - HTML"](#) display to view specific metrics about process behavior over a specified period of time and determine which activity in the process may be causing the bottleneck.



### Available Metrics

Select the metric driving the heatmap display. The default is **Alert Severity**. Each **Metric** has a color gradient bar that maps values to colors. The heatmap organizes the servers by host, where each rectangle represents a process. Mouse-over any rectangle to display the current values of the metrics for the process. Click on a rectangle to drill-down to the associated ["TIBCO BusinessWorks 5 Process Summary - HTML"](#) display for a detailed view of metrics for that particular process.

<b>Alert Severity</b>	<p>The maximum level of alerts in the heatmap rectangle. Values range from <b>0</b> - <b>2</b>, as indicated in the color gradient  bar, where <b>2</b> is the highest Alert Severity:</p> <ul style="list-style-type: none"> <li> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Alert Count</b>	<p>The total number of critical and warning alerts in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.</p>
<b>CPU Used %</b>	<p>The percent (%) CPU used in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the alert threshold of <b>BwProcessTotalCpuPercentHigh</b>, which is <b>100</b>. The middle value in the gradient bar indicates the middle value of the range (the default is <b>50</b>).</p>
<b>Completed</b>	<p>The total number of completed processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of completed processes in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>Active</b>	<p>The total number of active processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of active processes in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>Aborted</b>	<p>The total number of aborted processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of aborted processes in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>Suspended</b>	<p>The total number of suspended processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of suspended processes in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>Exec Time / s</b>	<p>The number of processes executed per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum execution rate of processes in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>Created / s</b>	<p>The number of processes created per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum created rate of processes in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>Aborted / s</b>	<p>The number of aborted processes per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum aborted rate of processes in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>Suspended / s</b>	<p>The number of suspended processes per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum suspended rate of processes in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>

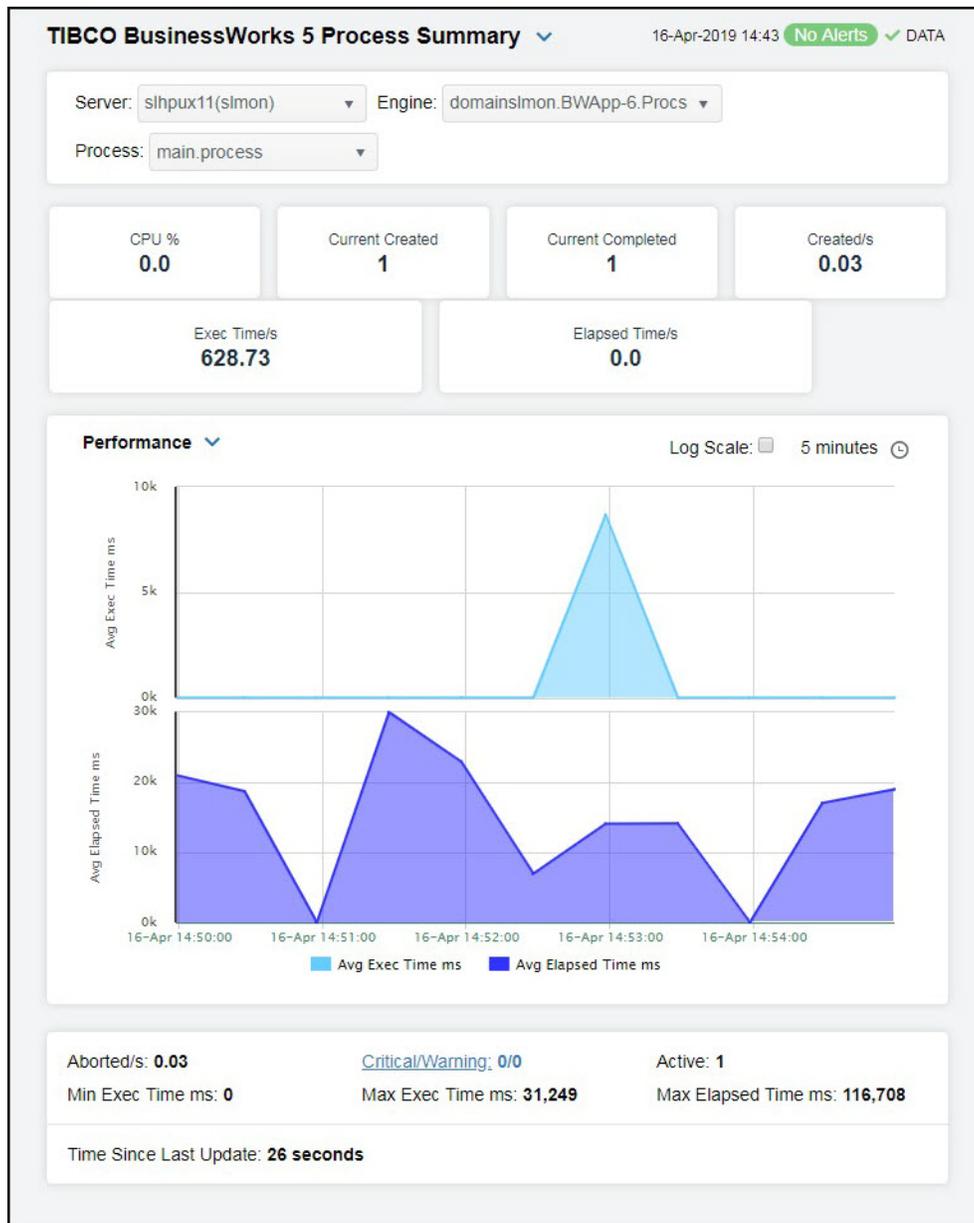
<b>Most Recent Exec Time</b>	The execution time for the most recently executed process in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the most recent execution time of processes in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
<b>Average Exec Time</b>	The average execution time for all processes in the heatmap rectangle, calculated by dividing the delta execution time for the interval by the delta completed, or the number of process instances that completed in the interval. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the average execution time of processes in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
<b>Most Recent Elapsed Time</b>	The elapsed time for the most recently executed process in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the most recent elapsed time of processes in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
<b>Average Elapsed Time</b>	The average elapsed time for all processes in the heatmap rectangle, calculated by dividing the delta elapsed time for the interval by the delta completed, or the number of process instances that completed in the interval. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the average elapsed time of processes in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

## TIBCO BusinessWorks 5 Process Summary - HTML

Clicking **BW5 Process Summary** in the left/navigation menu opens the **TIBCO BusinessWorks 5 Process Summary** display, which allows you to track utilization and performance metrics for specific BW5 processes. You can select a server, engine, and process from the drop-down menus. Clicking on the information boxes at the top of the display takes you to the "[TIBCO BusinessWorks 5 Processes Table - HTML](#)" display, where you can view additional data on processes.

There are two options in the trend graph: **Utilization** and **Performance**. In the **Utilization** option on the trend graph, you can view trend data for the rate of created processes and CPU percentage over a selected time range. In the **Performance** option on the trend graph, you can view trend data for used average execution time and average elapsed time over a selected time range.

Clicking the **Critical/Warning** link at the bottom of the display opens the Alerts Table by Component display.



## BW5 Activities - HTML

These displays present performance metrics and alert status for all BW5 activities. Clicking **BW5 Activities** from the left/navigation menu opens the “[TIBCO BusinessWorks 5 Activities Table - HTML](#)” display, where each row in the table displays all available metrics from the Hawk microagent for an activity. The options available under **BW5 Activities** are:

- **All Activities Heatmap:** Opens the “[TIBCO BusinessWorks 5 Activities Heatmap - HTML](#)”, which shows process execution metrics for all activities.
- **BW5 Activity Summary:** Opens the “[TIBCO BusinessWorks 5 Activity Summary - HTML](#)”

display, which shows historical and current performance metrics for a single activity, including average execution times and execution counts.

## TIBCO BusinessWorks 5 Activities Table - HTML

Select a server, engine, and process from the drop-down menus to see activities for the selected combination. Each row in the table is a different activity. Each table row displays all metrics available from the Hawk microagent for an activity. (Refer to documentation for TIBCO BusinessWorks Administration, see Appendix A: TIBCO Hawk Microagent Methods).

Click a column header to sort column data in ascending or descending order. Double-click on a table row to drill-down to the ["TIBCO BusinessWorks 5 Activity Summary - HTML"](#) display and view metrics for that particular activity. Toggle between the commonly accessed **Table** and **Heatmap** displays by clicking the drop down list on the display title.

An EXPIRED activity and the associated engine are deleted from displays when the associated server exceeds its specified threshold. Processes associated with the engine are also deleted from displays.

**TIBCO BusinessWorks 5 Activities Table** 16-Apr-2019 14:58 No Alerts DATA

Server:  Engine:   
 Process:

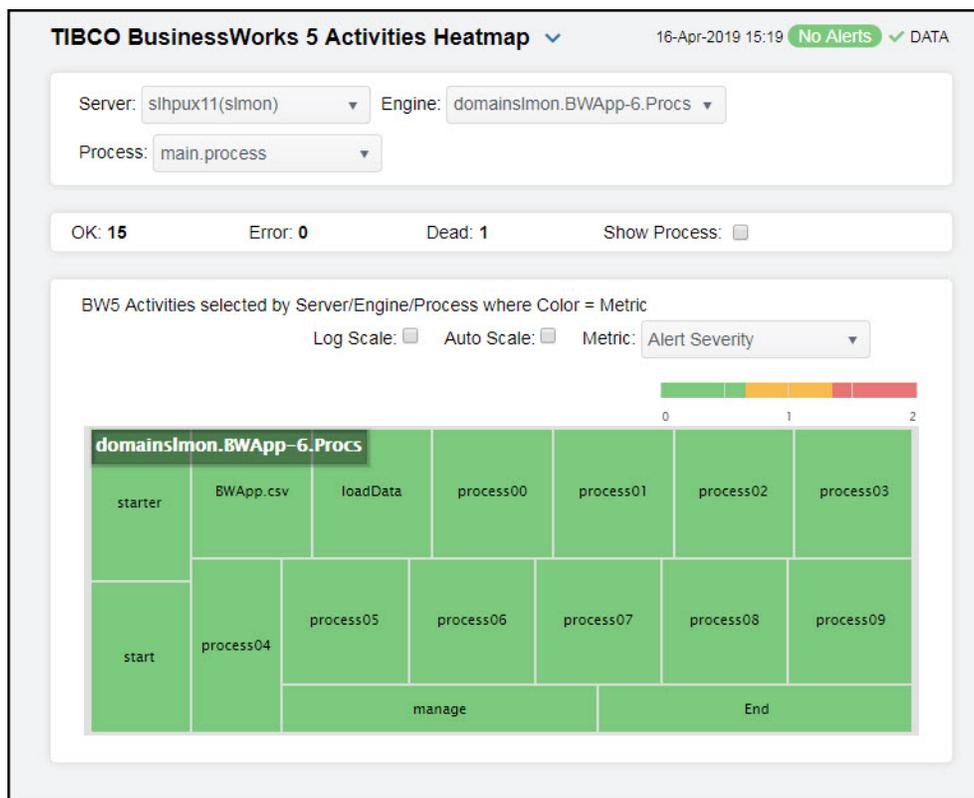
OK: **15**      Error: **0**      Dead: **1**

Activity	Alert Level	Alert Count	Time Since Last Update	Last Return Code
starter	✓	0	8 seconds	OK
start	✓	0	8 seconds	OK
BWApp.csv	✓	0	1 yr 21 days	OK
loadData	✓	0	1 yr 21 days	OK
process00	✓	0	8 seconds	OK
process01	✓	0	8 seconds	OK
process02	✓	0	8 seconds	OK
process03	✓	0	8 seconds	OK
process04	✓	0	8 seconds	OK
process05	✓	0	8 seconds	OK
process06	✓	0	8 seconds	OK
process07	✓	0	8 seconds	OK
process08	✓	0	8 seconds	OK

## TIBCO BusinessWorks 5 Activities Heatmap - HTML

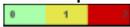
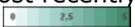
Clicking **All Activities Heatmap** in the left/navigation menu opens the **TIBCO BW5 Activities Heatmap**, which allows you to view the status and alerts of the execution times for all activities on all engines, or you can filter to look at specific servers, engines or processes. Use the **Metric** drop-down menu to view processes in the heatmap by the **Alert Severity, Alert Count, Executions, Errors, Execution Time per second, Error rate, Most Recent Execution Time**, and **Maximum Execution Time**.

The heatmap is organized by host with each rectangle representing an activity. Move your mouse over a node to display current metrics. Click on a node to drill-down to the ["TIBCO BusinessWorks 5 Activity Summary - HTML"](#) display to view specific metrics about process behavior over a specified period of time and determine which activity in the process may be causing the bottleneck.



### Available Metrics

Select the metric driving the heatmap display. The default is **Alert Severity**. Each **Metric** has a color gradient bar that maps values to colors. The heatmap organizes the servers by host, where each rectangle represents an activity. Mouse-over any rectangle to display the current values of the metrics for the activity. Click on a rectangle to drill-down to the associated ["TIBCO BusinessWorks 5 Activity Summary - HTML"](#) display for a detailed view of metrics for that particular activity.

<b>Alert Severity</b>	<p>The maximum level of alerts in the heatmap rectangle. Values range from <b>0</b> - <b>2</b>, as indicated in the color gradient  bar, where <b>2</b> is the highest Alert Severity:</p> <ul style="list-style-type: none"> <li> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</li> <li> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</li> <li> Green indicates that no metrics have exceeded their alert thresholds.</li> </ul>
<b>Alert Count</b>	<p>The total number of critical and warning alerts in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.</p>
<b>Executions</b>	<p>The total number of executed processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the number of executions in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>Errors</b>	<p>The total number of errors in the heatmap rectangle. The color gradient  bar populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the number of errors in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>Exec Time / sec</b>	<p>The number of processes executed per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the rate of executions in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>Errors / sec</b>	<p>The number of errors per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the rate of errors in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>Most Recent Exec Time</b>	<p>The execution time for the most recently executed process in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the most recent execution time in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>Max Exec Time</b>	<p>The maximum execution time for executed processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum execution time in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>

## TIBCO BusinessWorks 5 Activity Summary - HTML

Clicking **BW5 Activity Summary** in the left/navigation menu opens the **TIBCO BusinessWorks 5 Activity Summary** display, which allows you to track utilization and performance metrics for specific BW5 activities. You can select a server, engine, process, and activity from the drop-down menus. Clicking on the information boxes at the top of the display takes you to the "[TIBCO BusinessWorks 5 Activities Table - HTML](#)" display, where you can view additional data on activities.

There are two options in the trend graph: **Performance** and **Success Rate and Average Failures**. In the **Performance** option on the trend graph, you can view trend data for the average elapsed time and average execution time over a selected time range. In the **Success Rate and Average Failures** option on the trend graph, you can view trend data for used execution rate and average errors over a selected time range.

Clicking the **Critical/Warning** link at the bottom of the display opens the Alerts Table by Component display.

### TIBCO BusinessWorks 5 Activity Summary ▼ 16-Apr-2019 15:35 No Alerts DATA

Server:  Engine:   
Process:  Activity:

Errors <b>0</b>	Last Return Code <b>OK</b>	Executions/s <b>0.0</b>	Avg Execution Time ms <b>0</b>
Avg Elapsed Time ms <b>0</b>		Errors/s <b>0.0</b>	

#### Performance ▼ Log Scale: 5 minutes ⌵

Called Process Defs: **process01.process** [Critical/Warning: 0/0](#)

Domain: <b>domainslmon</b>	Current Exec Time ms: <b>0</b>	Exec Time ms: <b>8,486,394</b>
Max Exec Time ms: <b>4,933</b>	Current Elapsed Time ms: <b>0</b>	Elapsed Time ms: <b>8,486,394</b>
Max Elapsed Time ms: <b>4,933</b>		

Time Since Last Update: **14 seconds**

---

## TIBCO Enterprise Message Service

The following TIBCO Enterprise Message Service Views (and their associated displays) can be found under **Components** tab > **Middleware**. The displays within the Views will be populated with data once the Solution Package for TIBCO Enterprise Message Service is configured in the RTView DataServer for TIBCO and the RTView DataServer for TIBCO is connected to RTViewCentral.

- "All EMS Servers"
- "Single EMS Server"
- "EMS Topics"
- "EMS Queues"
- "EMS Clients"

### All EMS Servers

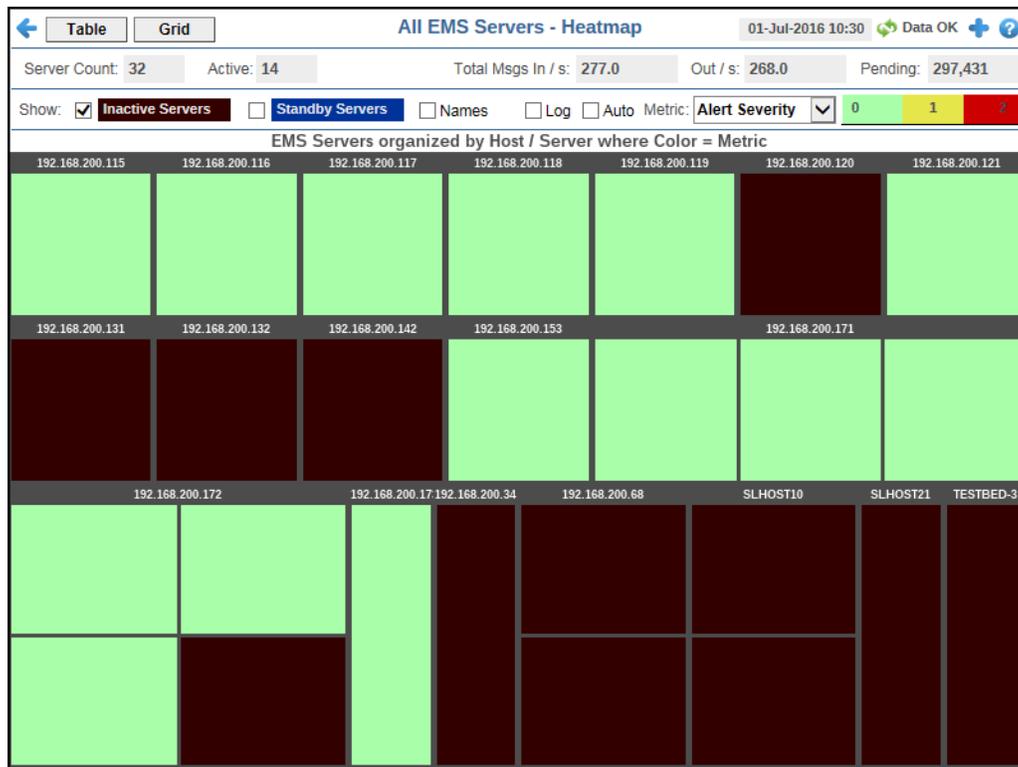
These displays present performance metrics and alert status for all EMS servers. The first three displays show different views of the same data:

- "All Servers Heatmap": Heatmap shows server and alert status for all EMS servers.
- "All Servers Table": Table shows all available utilization metrics for all EMS servers.
- "All Servers Grid": Grid enables you to see general performance of EMS servers in parallel. If you have few servers, this display is useful for verifying servers are active and generally performing as expected.
- "All Servers Topology": Topology of server routes and connections, as well as the status of active servers and standby servers that form a fault-tolerant pair.

### All Servers Heatmap

View status and alerts of all EMS servers. Use the **Metric** drop-down menu to view the **Alert Severity, Alert Count, Connections, Pending Messages, Inbound Message Rate, Outbound Message Rate, or Message Memory Percent (%)**.

The heatmap is organized by host, each rectangle representing a server. The rectangle color indicates the most critical alert state. Click on a node to drill-down to the Single Server Summary display and view metrics for a particular server. Toggle between the commonly accessed **Table, Grid, and Heatmap** displays. Mouse-over rectangles to view more details about host performance and status.



**Title Bar** (possible features are):

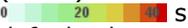
- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- The number of items currently in the display.
- Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

**Fields and Data**

This display includes:

- Server Count** The total number of active, inactive, and standby EMS servers.
- Active** The total number of currently active EMS servers.
- Total Msgs In/s** **In/s** The total number of inbound messages, per second, from all producers and consumers on all EMS servers.
- Out/s** The total number of outbound messages, per second, from all producers and consumers on all EMS servers.
- Pending** The total number of pending messages waiting to be processed on all EMS servers. Click to open the ["All Servers Table"](#) display.

<b>Show</b>	Select the type of servers for which to display data. By default, all active servers are displayed.
<b>Inactive Servers</b>	Select to include servers that are not currently running. <b>Inactive Servers</b> are represented in dark red.
<b>Standby Servers</b>	Select to include servers that are currently in Standby mode. Standby Servers are represented in blue.
<b>Names</b>	Select to display the names of servers on the hosts.
<b>Log</b>	This option enables visualization on a logarithmic scale, and should be used when the range in your data is very broad. For example, if you have data that ranges from the tens to the thousands, then data in the range of tens will be neglected visually if you do not check this option. This option makes data on both extreme ranges visible by using the logarithmic of the values rather than the actual values.
<b>Auto</b>	When checked, the values of the selected metric are auto-scaled to its highest defined value. When unchecked, the values of the selected metric display based on the threshold defined for the alert associated with the selected metric. Selecting Auto helps to visualize the range of the values currently present for the selected metric instead of the threshold of the alert that has been associated with the metric. All metrics that have not been associated in the heatmap defaults with alerts use a monochromatic color gradient bar (whites and greens). All metrics that have been associated in the heatmap defaults with alerts use a multi-chromatic color gradient bar (reds, yellows, white, and greens).
<b>Metric</b>	Select the metric driving the heatmap display. The default is Alert Severity. Each <b>Metric</b> has a color gradient bar that maps values to colors. The heatmap organizes the servers by host, where each rectangle represents a server. Mouse-over any rectangle to display the current values of the metrics for the Server. Click on a rectangle to drill-down to the associated <a href="#">"Single Server Summary"</a> display for a detailed view of metrics for that particular server.
<b>Alert Severity</b>	<p>The maximum alert level in the item (index) associated with the rectangle. Values range from <b>0</b> to <b>2</b>, as indicated in the color gradient bar , where <b>2</b> is the greatest <b>Alert Severity</b>.</p> <p><b>2</b> -- Metrics that have exceeded their specified <b>ALARMLEVEL</b> threshold and have an Alert Severity value of <b>2</b> are shown in red. For a given rectangle, this indicates that one or more metrics have exceeded their alarm threshold.</p> <p><b>1</b> -- Metrics that have exceeded their specified <b>WARNINGLEVEL</b> threshold and have an Alert Severity value of <b>1</b> are shown in yellow. For a given rectangle, this indicates that one or more metrics have exceeded their warning threshold.</p> <p><b>0</b> -- Metrics that have not exceeded either specified threshold have an Alert Severity value of <b>0</b> and are shown in green. For a given rectangle, this indicates that no metrics have exceeded a specified alert threshold.</p>
<b>Alert Count</b>	<p>The total number of alarm and warning alerts in a given item (index) associated with the rectangle.</p> <p>The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>

<b>Connections</b>	<p>The total number of connections in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of connections in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p> <p>The <b>Auto</b> option does not impact this metric.</p>
<b>Pend Messages</b>	<p>The total number of pending messages in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from <b>0</b> to the alert threshold of <b>EmsServerPendingMsgsHigh</b>, which is <b>3500</b>. The middle value in the gradient bar indicates the middle value of the range (the default is <b>1750</b>).</p> <p>When <b>Auto</b> is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>
<b>In Msg Rate</b>	<p>The total number of inbound messages in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from <b>0</b> to the alert threshold of <b>EmsServerInMsgRateHigh</b>, which is <b>40</b>. The middle value in the gradient bar indicates the middle value of the range (the default is <b>20</b>).</p> <p>When <b>Auto</b> is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>
<b>Out Msg Rate</b>	<p>The total number of outbound messages in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from <b>0</b> to the alert threshold of <b>EmsServerOutMsgRateHigh</b>, which is <b>40</b>. The middle value in the gradient bar indicates the middle value of the range (the default is <b>20</b>).</p> <p>When <b>Auto</b> is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>
<b>Mem Msg %</b>	<p>The percent (%) memory used by messages in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from <b>0</b> to the alert threshold of <b>EmsServerMemUsedHigh</b>, which is <b>40</b>. The middle value in the gradient bar indicates the middle value of the range (the default is <b>20</b>).</p> <p>When <b>Auto</b> is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>

## All Servers Table

Investigate detailed utilization metrics for all EMS servers. The **All Servers Table** contains all metrics available for servers, including the number of current client connections. Each row in the table contains data for a particular server. Click a column header to sort column data in numerical or alphabetical order. Click on a table row to drill-down to the “[Single Server Summary](#)” display and view metrics for that particular server. Toggle between the commonly accessed **Table**, **Grid**, and **Heatmap** displays.

Heatmap Grid EMS All Servers Table 01-Jul-2016 10:37 Data OK

Server Count: 32 Active: 14 Total Msgs In / s: 171.0 Out / s: 211.0 Pending: 297,459

Show:  Inactive Servers  Standby Servers

URL	serverName	Host	Expired	Alert Level	state	versionIn
tcp://192.168.200.115:7222	EMS-SERVER	192.168.200.115	<input type="checkbox"/>		Active	6.0.0.8
tcp://192.168.200.116:7222	EMS-SERVER	192.168.200.116	<input type="checkbox"/>		Active	6.0.0.8
tcp://192.168.200.117:7222	EMS-SERVER	192.168.200.117	<input type="checkbox"/>		Active	6.1.0.6
tcp://192.168.200.118:7222	EMS-SERVER	192.168.200.118	<input type="checkbox"/>		Active	6.3.0.5
tcp://192.168.200.119:7222	EMS-SERVER	192.168.200.119	<input type="checkbox"/>		Active	6.3.0.5
tcp://192.168.200.120:7222	Unknown (tcp://192.168.200...	192.168.200.120	<input checked="" type="checkbox"/>			
tcp://192.168.200.121:7222	EMS-SERVER	192.168.200.121	<input type="checkbox"/>		Active	6.0.0.8
tcp://192.168.200.131:7222	TESTBED-1 (tcp://192.168.2...	192.168.200.131	<input checked="" type="checkbox"/>			
tcp://192.168.200.132:7222	TESTBED-2 (tcp://192.168.2...	192.168.200.132	<input checked="" type="checkbox"/>			
tcp://192.168.200.142:7222	Unknown (tcp://192.168.200...	192.168.200.142	<input checked="" type="checkbox"/>			
tcp://192.168.200.153:7222	EMS-SERVER-153	192.168.200.153	<input type="checkbox"/>		Active	8.2.2.3
tcp://192.168.200.171:6010	EMS-SLDEMOS2-6010	192.168.200.171	<input type="checkbox"/>		Active	6.0.0.8
tcp://192.168.200.171:6020	EMS-SLDEMOS2-6020	192.168.200.171	<input type="checkbox"/>		Active	6.0.0.8
tcp://192.168.200.171:6030	EMS-SLDEMOS2-6030	192.168.200.171	<input type="checkbox"/>		Active	6.0.0.8
tcp://192.168.200.172:8011	EMS-SLDEMOS3-8010	192.168.200.172	<input type="checkbox"/>		Active	6.0.0.8
tcp://192.168.200.172:8020	EMS-SLDEMOS3-8020	192.168.200.172	<input type="checkbox"/>		Active	6.0.0.8
tcp://192.168.200.172:8030	EMS-SLDEMOS3-8030	192.168.200.172	<input type="checkbox"/>		Active	6.0.0.8
tcp://192.168.200.172:8031	Unknown (tcp://192.168.200...	192.168.200.172	<input checked="" type="checkbox"/>			
tcp://192.168.200.173:9010	EMS-SLDEMOS4-9010	192.168.200.173	<input type="checkbox"/>		Active	6.0.0.8
tcp://192.168.200.34:7222	TESTBED-34 (tcp://192.168...	192.168.200.34	<input checked="" type="checkbox"/>			
tcp://192.168.200.68:7222	Unknown (tcp://192.168.200...	192.168.200.68	<input checked="" type="checkbox"/>			
tcp://192.168.200.68:7224	Unknown (tcp://192.168.200...	192.168.200.68	<input checked="" type="checkbox"/>			
tcp://SLHOST10:7010	Unknown (tcp://SLHOST10:7...	SLHOST10	<input checked="" type="checkbox"/>			
tcp://SLHOST10:7011	Unknown (tcp://SLHOST10:7...	SLHOST10	<input checked="" type="checkbox"/>			
tcp://SLHOST21:7222	EMS-SERVER-SLHOST21 (t...	SLHOST21	<input checked="" type="checkbox"/>			
tcp://TESTBED-3:7022	Unknown (tcp://TESTBED-3:...	TESTBED-3	<input checked="" type="checkbox"/>			

**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu Table open commonly accessed displays.
- 6,047 The number of items currently in the display.
- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

**Fields and Data**

This display includes:

- Server Count** The total number of active, inactive and standby EMS servers. **Inactive Servers** are represented in dark red. **Standby Servers** are represented in blue.
- Active** The total number of currently active EMS servers.
- Total Msgs In/s** The total number of inbound messages, per second, from all producers and consumers on all EMS servers.
- Out/s** The total number of outbound messages, per second, from all producers and consumers on all EMS servers.
- Pending** The total number of inbound and outbound messages waiting to be processed on all EMS servers.

<b>Show</b>	Select the type of servers to display data for. By default, all active servers are displayed.
<b>Inactive Servers</b>	Select to include servers that are not processing requests in the table. <b>Inactive Servers</b> are represented in dark red.
<b>Standby Servers</b>	Select to include servers that are not currently running. <b>Standby Servers</b> are represented in blue.
<b>Table</b>	This table shows information for all EMS servers. Click on a table row to drill-down to the <a href="#">"Single Server Summary"</a> display and view metrics for that particular server.
<b>URL</b>	Select to include servers that are currently in Standby mode. <b>Standby Servers</b> are represented in blue.
<b>serverName</b>	The name of the server.
<b>Host</b>	The name or IP address for the host server.
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > <b>(Project Name)</b> > <b>Solution Package Configuration</b> > <b>TIBCO Enterprise Message Service</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Alert Level</b>	<p>The maximum alert level in the item (index) associated with the rectangle. Values range from 0 to 2, as indicated in the color gradient bar, where <b>2</b> is the greatest Alert Severity.</p> <p> -- One or more alerts have exceeded their specified <b>ALARMLEVEL</b> threshold, have an Alert Severity value of <b>2</b>, and are shown in red.</p> <p> -- One or more alerts have exceeded their specified <b>WARNINGLEVEL</b> threshold, have an Alert Severity value of <b>1</b>, and are shown in yellow.</p> <p> -- No alerts have exceeded an alert threshold, which have an Alert Severity value of <b>0</b>, and are shown in green.</p>
<b>state</b>	<p>The server status:</p> <p><b>Active</b> -- The server is currently processing requests.</p> <p><b>Inactive</b> -- The server is not currently processing requests. <b>Inactive Servers</b> are represented in dark red.</p> <p><b>Standby</b> -- The server is functioning as a backup for a primary server. <b>Standby Servers</b> are represented in blue.</p>
<b>versionInfo</b>	The TIBCO EMS software version currently running.
<b>faultTolerantURL</b>	The IP address and port number for the source (application, server, and so forth) associated with the alert.
<b>asyncDBsize</b>	The amount of database space, in bytes, occupied by asynchronous data on the server.
<b>backupName</b>	The name of the backup server assigned as the backup to this server.
<b>connectionCount</b>	The number of clients currently connected to the server.
<b>diskReadRate</b>	The speed at which the server reads disk data.
<b>diskWriteRate</b>	The speed at which the server writes data to disk.

<b> durableCount </b>	The number of durables on the server.
<b> inboundBytesRate </b>	The rate of inbound messages in bytes per second.
<b> inboundMessageCount </b>	The number of inbound messages received by the server since the server was started.
<b> inboundMessageRate </b>	The rate of inbound messages in number of messages per second.
<b> MaxMessageMemory </b>	The maximum amount of memory, in bytes, allocated for use by messages on the server.
<b> messageMemory </b>	The amount of memory, in bytes, currently used by messages on the server.
<b> messageMemoryPct </b>	The amount of memory, in percent, used by messages on the server.
<b> messageMemoryPooled </b>	The currently allocated pool size, in bytes, for messages.
<b> outboundBytesRate </b>	The rate of outbound messages in bytes per second.
<b> outboundMessageCount </b>	The number of outbound messages sent by the server since the server was started.
<b> outboundMessageRate </b>	The rate of outbound messages in number of messages per second.
<b> pendingMessageCount </b>	The number of currently pending messages on the server.
<b> pendingMessageSize </b>	The amount of space, in bytes, pending messages use on the server.
<b> processId </b>	The process ID of the EMS server.
<b> queueCount </b>	The number of message queues.
<b> startTime </b>	The date and time that the server was started.
<b> syncDBSize </b>	The amount of database space, in bytes, occupied by synchronous data on the server.
<b> topicCount </b>	The number of currently active topics on the server.
<b> upTime </b>	The amount of time, in milliseconds, since the server was started.
<b> time_stamp </b>	The date and time this row of data was last updated.

## All Servers Grid

Track and view in parallel the general performance of all EMS servers. Click on a node to drill-down to the "Single Server Summary" display and view detailed metrics for that particular server.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** **Table** open commonly accessed displays.
- 6,047** The number of items currently in the display.

- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

### Fields and Data

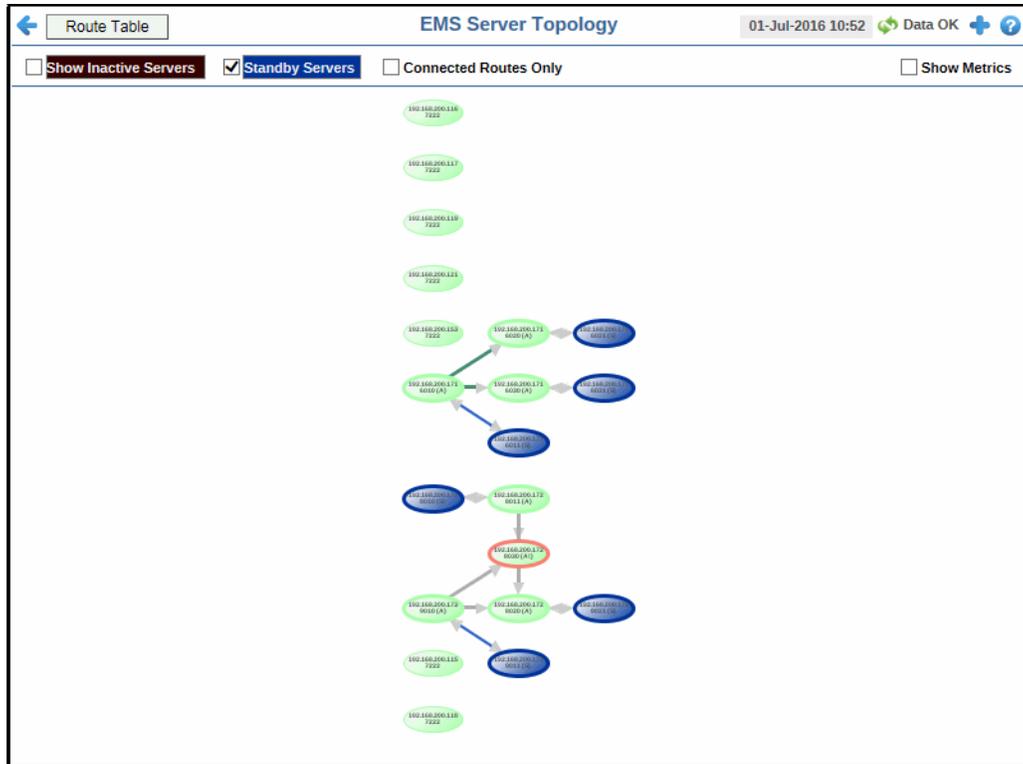
This display includes:

- Server Count** The total number of active, inactive and standby EMS servers. Inactive Servers are represented in dark red. Standby Servers are represented in blue.
- Active** The total number of currently active EMS servers.

<b>Total Msgs</b>	<b>In/s</b>	The total number of inbound messages, per second, from all producers and consumers on all EMS servers.
	<b>Out/s</b>	The total number of outbound messages, per second, from all producers and consumers on all EMS servers.
	<b>Pending</b>	The total number of inbound and outbound messages waiting to be processed on all EMS servers. Click to open the <a href="#">“All Servers Table”</a> display.
<b>Show</b>		Select the type of servers to display data for. By default, all active servers are displayed.
	<b>Inactive Servers</b>	Select to include servers that are not processing requests in the table. <b>Inactive Servers</b> are represented in dark red.
	<b>Standby Servers</b>	Select to include servers that are not currently running. <b>Standby Servers</b> are represented in blue.
<b>Sort By</b>	<b>Server Name</b>	Select to organize the servers in the grid by server name.
	<b>Server URL</b>	Select to organize the servers in the grid by server URL.
<b>Descending</b>		When checked, lists servers in the grid in descending order.
<b>Time Range</b>		Select a time range from the drop down menu varying from <b>2 Minutes to Last 7 Days</b> , or display <b>All Data</b> .
<b>Grid</b>	<b>Server Name</b>	The name of the server.
	<b>URL</b>	The URL for the server.
	<b>Uptime</b>	The amount of time, in milliseconds, since the server was started.
	<b>Pend Msgs</b>	The number of currently pending messages on the server.
	<b>State</b>	The server status: <b>Active</b> -- The server is currently processing requests. <b>Inactive</b> -- The server is not currently processing requests. Inactive Servers are represented in dark red. <b>Standby</b> -- The server is functioning as a backup for a primary server. Standby Servers are represented in blue.
	<b>In Rate</b>	The rate of inbound messages in messages per second.
	<b>Out Rate</b>	The rate of outbound messages in messages per second.
	<b>Trend Graphs</b>	Shows message data for the server. <b>Pend</b> -- Traces the total number of pending messages on the server. <b>In</b> -- Traces the rate of inbound messages in messages per second. <b>Out</b> -- Traces the rate of outbound messages in messages per second.

## All Servers Topology

View a server topology map for all EMS servers. Click on a node to drill-down to the “[Single Server Summary](#)” display and view metrics for that particular server.



**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- The number of items currently in the display.

- Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

**Note:** Clicking the **Route Table** button displays the **EMS Server Route Table** window. See “[EMS Server Route Table](#)” for more information.

### Fields and Data

This display includes:

<b>Show</b>	The total number of active, inactive and standby EMS servers. Inactive Servers are represented in dark red. Standby Servers are represented in blue.
<b>Inactive Servers</b>	Select to show servers that are not processing requests in the topology. Inactive Servers are represented in dark red.
<b>Standby Servers</b>	Select to show servers that are not processing requests in the topology. Standby Servers are represented in blue.
<b>Connected Routes Only</b>	Select to show only routes that have an active connection.
<b>Show Metrics</b>	Available on desktop application deployments only. Shows the total input message rates, per second, on the top of each server icon and the total output message rate on the bottom of each server icon.
<b>Topology</b>	<p>Routes are shown between the active server and the standby server, which form a fault-tolerant pair. Either of the servers in a fault-tolerant pair can become the active server or the standby server. <b>Show Standby Servers</b> and <b>Show Inactive Servers</b> enable you to include or exclude standby and inactive servers. <b>Inactive Servers</b> are represented in dark red. <b>Standby Servers</b> are represented in blue. By default, standby servers are included in the topology and inactive servers are not.</p> <p>Typically, it takes about 30 seconds for a server to appear in the display after startup.</p> <p>The active server in a fault-tolerant pair appears in green with the suffix <b>(A)</b> appended to its URL. The standby server appears in blue, with the suffix <b>(S)</b> appended to its URL. Their link is blue and labeled <b>FT</b>.</p> <p>If the active server fails:</p> <ul style="list-style-type: none"> <li>the failed server becomes inactive, its suffix changes to <b>(X!)</b>, and the node turns red with a red outline.</li> <li>the standby server becomes active, its suffix changes to <b>(A!)</b>, and the node turns green with a red outline.</li> <li>the link between the two servers turns red.</li> </ul> <p>If the standby server fails:</p> <ul style="list-style-type: none"> <li>the failed server becomes inactive, its suffix changes to <b>(X!)</b>, and the node turns red with a red outline.</li> <li>the active servers' suffix changes to <b>(A!)</b> and it is outlined in red.</li> <li>the link between the two servers turns red.</li> </ul> <p>If a failed server recovers:</p> <ul style="list-style-type: none"> <li>the recovered server becomes the standby server, its suffix changes to <b>(S)</b>, and the node turns blue with a grey outline.</li> <li>the active servers' suffix <b>(A!)</b> changes to <b>(A)</b>, and the red node outline changes back to grey.</li> <li>the link between the two servers changes back to blue.</li> </ul> <p><b>Suffix Definition</b></p> <p><b>A</b> -- This is the active server and it is running.  <b>A!</b> -- This is the active server and it is running but its standby has failed.  <b>S</b> -- This is the standby server and it is running.  <b>X!</b> -- The server is inactive.</p> <p><b>Node Color Definition</b></p> <p>-- This is the active server and it is running.  <b>Blue</b> -- This is the standby server and it is in standby mode.  -- The server is inactive.</p>

- Link Color Definition**      **Blue** -- The two servers in the pair are running.  
 -- One of the servers in the pair is inactive.
- Outline Color Definition**      **Grey** -- The two servers in the pair are running.  
 -- One of the servers in the pair is inactive. If the node color indicates this server is running, its pair is inactive.

### EMS Server Route Table

Displays metrics for server routes on all servers. Inbound metrics, such as **inboundByteRate**, indicate an in route to the server. Outbound metrics, such as **outboundByteRate**, indicate an out route to the server.

EMS Server Route Table						01-Jul-2016 10:59	Data OK	+	?
remoteURL	remoteName	connected	stalled	inboundByteRate	inboundMessageRate				
tcp://192.168.200.171:6020,tcp://192.168.200.1...	EMS-SLDEMOS2-6020		0	0.0					
tcp://192.168.200.171:6020,tcp://192.168.200.1...	EMS-SLDEMOS2-6020		0	0.0					
tcp://192.168.200.171:6030,tcp://192.168.200.1...	EMS-SLDEMOS2-6030		0	0.0					
tcp://192.168.200.171:6030,tcp://192.168.200.1...	EMS-SLDEMOS2-6030		0	0.0					
tcp://192.168.200.172:8020,tcp://192.168.200.1...	EMS-SLDEMOS3-8020		0	0.0					
tcp://192.168.200.172:8020,tcp://192.168.200.1...	EMS-SLDEMOS3-8020		0	0.0					
tcp://192.168.200.172:8020,tcp://192.168.200.1...	EMS-SLDEMOS3-8020		0	0.0					
tcp://192.168.200.172:8020,tcp://192.168.200.1...	EMS-SLDEMOS3-8020		0	0.0					
tcp://192.168.200.172:8030,tcp://192.168.200.1...	EMS-SLDEMOS3-8030		0	0.0					
tcp://192.168.200.172:8030,tcp://192.168.200.1...	EMS-SLDEMOS3-8030		0	0.0					
tcp://192.168.200.172:8030,tcp://192.168.200.1...	EMS-SLDEMOS3-8030		0	0.0					
tcp://192.168.200.172:8030,tcp://192.168.200.1...	EMS-SLDEMOS3-8030		0	0.0					
tcp://localhost:7022	EMS-SERVER2		0	0.0					
tcp://localhost:7022	EMS-SERVER2		0	0.0					
tcp://localhost:7022	EMS-SERVER2		0	0.0					
tcp://localhost:7022	EMS-SERVER2		0	0.0					
tcp://SLHOST10	EMS-SLDEMOS1-7010		0	0.0					
tcp://vmrh5-4	EMS-SLDEMOS2-6010		0	0.0					
tcp://vmrh5-4	EMS-SLDEMOS2-6010		0	0.0					

#### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- The number of items currently in the display.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

Open the Alert Views - RTView Alerts Table display.

### Fields and Data

This display includes:

- remoteURL**      The remote URL of the server.
- remoteName**      The name of the server.
- connected**      The connection state of the server route.

	<ul style="list-style-type: none"> <li> -- One or more routes for this server are disconnected.</li> <li> -- All routes for this server are connected.</li> <li> -- There are no routes for this server.</li> </ul>
<b>stalled</b>	<p>Indicates whether the IO flow stalled on the route.</p> <p>A value of <b>0</b> (zero) = not stalled.</p> <p>A value of <b>1</b> = stalled.</p>
<b>inboundByteRate</b>	The rate of inbound data in bytes, per second.
<b>inboundMessageRate</b>	The rate of inbound messages in number of messages per second.
<b>inboundTotalBytes</b>	The total number of inbound bytes.
<b>inboundTotalMessages</b>	The total number of inbound messages.
<b>outboundByteRate</b>	The rate of inbound data in bytes, per second.
<b>outboundMessageRate</b>	The rate of outbound messages in number of messages per second.
<b>outboundTotalBytes</b>	The total number of outbound bytes.
<b>outboundTotalMessages</b>	The total number of outbound messages.
<b>zoneName</b>	The name of the zone for the route.
<b>zoneType</b>	Indicates a multi-hop or one-hop route.
<b>active</b>	<p>Indicates whether the server route is currently transferring data:</p> <p><b>1</b> = true (is transferring data)</p> <p><b>0</b> = false</p>
<b>inactive</b>	<p>Indicates whether the server route is not currently transferring data:</p> <p><b>1</b> = true (is <b>not</b> transferring data)</p> <p><b>0</b> = false</p>
<b>suspended</b>	<p>Indicates whether outbound messages to the route have been suspended:</p> <p><b>1</b> = true</p> <p><b>0</b> = false</p>
<b>remoteURLName</b>	The IP address and name for the remote connection.

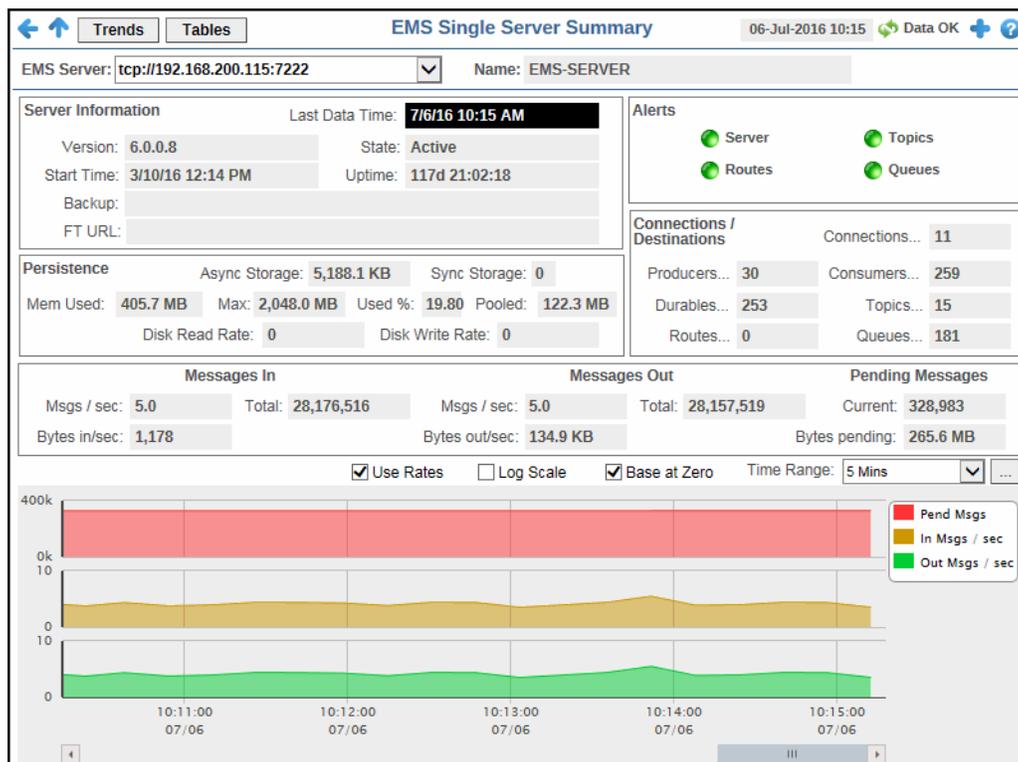
## Single EMS Server

These displays present detailed performance metrics, alert status and connection information for a single EMS server.

- **"Single Server Summary"**: Shows information for a single EMS server such as server connection details, the number of client connections, memory utilization, message performance metrics and alert status.
- **"Single Server Trends"**: Trend graphs show utilization metrics for a single EMS server, such as the number of client connections, number of pending messages and in/out rate, and memory and disk utilization.
- **"Single Server Tables"**: Tables show information about how the Monitor is connected to the EMS server, metrics queried from the server and alert details.

### Single Server Summary

Track utilization and performance metrics for specific servers.



#### Title Bar (possible features are):

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu, Table open commonly accessed displays.
- 6,047 The number of items currently in the display.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

Open the Alert Views - RTView Alerts Table display.

**Fields and Data**

This display includes:

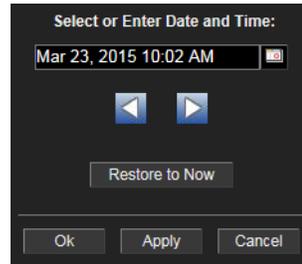
<b>EMS Server</b>	Select the EMS Server for which you want to view data. The selection made here populates this display.
<b>Name</b>	The name of the EMS Server selected from the EMS Server drop-down menu.
<b>Server Information</b>	<p><b>Version</b> The TIBCO EMS software version currently running.</p> <p><b>Start Time</b> The date and time that the server was started.</p> <p><b>Backup</b> The name of the backup server for the server.</p> <p><b>FT URL</b> The IP address and port number, or the hostname and port number, of the fault tolerant standby server assigned to this server.</p> <p><b>Last Data Time</b> The time that a data update was last made.</p> <p><b>State</b> The server status:  <b>Active</b> -- The server is currently processing requests.  <b>Inactive</b> -- The server is not currently processing requests.  <b>Standby</b> -- The server is functioning as a backup for a primary server.</p> <p><b>Uptime</b> The amount of time since the server was started.  Format:  <b>dd HH:MM:SS</b>  <b>&lt;days&gt; &lt;hours&gt;:&lt;minutes&gt;:&lt;seconds&gt;</b>  For example:  <b>10d 08:41:38</b></p>
<b>Persistence</b>	<p><b>Async Storage</b> The amount of database space, in bytes, used by asynchronous message persistence data on the server</p> <p><b>Sync Storage</b> The amount of database space, in bytes, used by synchronous message persistence data on the server.</p> <p><b>Mem Used</b> The amount of memory, in kilobytes, used by message persistence on the server.</p> <p><b>Max</b> The maximum amount of memory, in kilobytes, used by message persistence on the server.</p> <p><b>Used %</b> The amount of memory, in percent, used by message persistence.</p> <p><b>Pooled</b> The amount of message memory that has been pooled.</p> <p><b>Disk Read Rate</b> The speed at which the server reads message persistence disk data.</p> <p><b>Disk Write Rate</b> The speed at which the server writes message persistence disk data.</p>
<b>Alerts</b>	<p><b>Server</b> Status indicator for server-related alerts. Click to open the EMS <a href="#">"Single Server Tables"</a> display and view the <b>Server Alert Table</b> for more detail.</p> <p> -- No alerts have exceeded a specified threshold.</p> <p> -- One or more alerts have exceeded their specified <b>WARNINGLEVEL</b> threshold.</p> <p> -- One or more alerts have exceeded their specified <b>ALARMLEVEL</b> threshold.</p>

	<b>Routes</b>	Status indicator for route-related alerts. Click to open the EMS <a href="#">"Single Server Tables"</a> display and view the <b>Server Alert Table</b> for more detail.  -- No alerts have exceeded a specified threshold.  -- One or more alerts have exceeded their specified <b>WARNINGLEVEL</b> threshold.  -- One or more alerts have exceeded their specified <b>ALARMLEVEL</b> threshold.
	<b>Topics</b>	Status indicator for topic-related alerts. Click to open the EMS <a href="#">"Single Server Tables"</a> display and view the <b>Server Alert Table</b> for more detail.  -- No alerts have exceeded a specified threshold.  -- One or more alerts have exceeded their specified <b>WARNINGLEVEL</b> threshold.  -- One or more alerts have exceeded their specified <b>ALARMLEVEL</b> threshold.
	<b>Queues</b>	Status indicator for queue-related alerts. Click to open the EMS <a href="#">"Single Server Tables"</a> display and view the <b>Server Alert Table</b> for more detail.  -- No alerts have exceeded a specified threshold.  -- One or more alerts have exceeded their specified <b>WARNINGLEVEL</b> threshold.  -- One or more alerts have exceeded their specified <b>ALARMLEVEL</b> threshold.
<b>Connections / Destinations</b>		Shows connection information for the server. The counts shown here are also visible in the <a href="#">"EMS Topics"</a> and <a href="#">"EMS Clients"</a> displays.
	<b>Producers</b>	The number of producers currently active on the server. Click to open the <a href="#">"EMS Clients"</a> / <a href="#">"Producers"</a> for Server display for details.
	<b>Durables</b>	The number of durables currently active on the server. Click to open the <a href="#">"EMS Clients"</a> / <a href="#">"Consumer Summary"</a> for Server display for details.
	<b>Routes</b>	The number of routes defined on the server.
	<b>Connections</b>	The number of clients currently connected to the server. Click to open the <a href="#">"EMS Clients"</a> / <a href="#">"Connections"</a> for Server display for details.
	<b>Consumers</b>	The number of consumers currently connected to the server. Click to open the <a href="#">"EMS Clients"</a> / <a href="#">"Producer Summary"</a> for Server display for details.
	<b>Topics</b>	The number of topics currently active on the server. Click to open the <a href="#">"EMS Topics"</a> / <a href="#">"All Topics Table"</a> display for details.
	<b>Queues</b>	The number of queues currently active on the server. Click to open the <a href="#">"EMS Topics"</a> / <a href="#">"All Queues Heatmap"</a> display for details.
<b>Messages In</b>	<b>Msgs/sec</b>	The number of inbound messages, per second, from all producers and consumers
	<b>Bytes in/sec</b>	The total size of inbound messages, in bytes per second, from all producers and consumers.
	<b>Total</b>	The total number of inbound messages, in bytes, from all producers and consumers since the server was started.

<b>Messages Out</b>	<b>Msgs/sec</b>	The number of outbound messages, per second, from all producers and consumers.
	<b>Bytes out/sec</b>	The total size of outbound messages, in bytes per second, from all producers and consumers.
	<b>Total</b>	The total of outbound messages, in bytes, from all producers and consumers since the server was started.
<b>Pending Messages</b>	<b>Current</b>	The total number of inbound and outbound messages currently waiting to be processed.
	<b>Bytes pending</b>	The total size of inbound and outbound messages, in bytes, currently waiting to be processed.
<b>Trend Graphs</b>	Shows message metrics for the selected server.	
	<b>Pend Message</b>	-- Traces the total number of inbound and outbound messages currently waiting to be processed.
	<b>In Msgs / sec</b>	-- Traces the number of inbound messages, per second, from all producers and consumers. This trend graph only displays when <b>Use Rates</b> is selected.
	<b>Out Msgs / sec</b>	-- Traces the number of outbound messages, per second, from all producers and consumers. This trend graph only displays when <b>Use Rates</b> is selected.
	<b>Delta In Msgs</b>	-- Traces the change in total inbound messages since the last update. This trend graph only displays when <b>Use Rates</b> is not selected.
	<b>Delta Out Msgs</b>	-- Traces the change in total outbound messages since the last update. This trend graph only displays when <b>Use Rates</b> is not selected.
	<b>Use Rates</b>	When this check box is selected, the inbound and outbound message rates ( <b>In Msgs/sec</b> and <b>Out Msgs/sec</b> ) display in the trend graph. When this check box is not selected, the delta inbound and outbound messages ( <b>Delta In Msgs</b> and <b>Delta Out Msgs</b> ) display in the trend graph.
<b>Log Scale</b>	This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.	

**Base at Zero** When this option is checked, zero is set as the Y axis minimum for all graph traces.

**Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



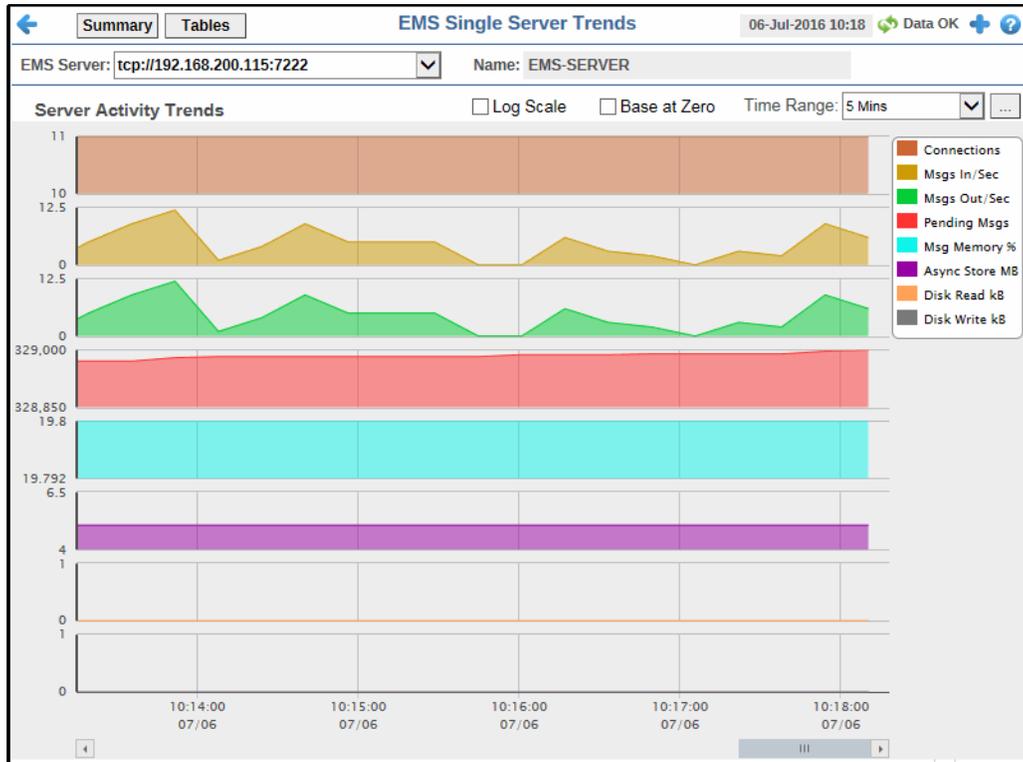
By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Single Server Trends

View trend graphs in parallel to investigate performance issues for a specific server.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** **Table** open commonly accessed displays.
- 6,047** The number of items currently in the display.

- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

### Fields and Data

This display includes:

- EMS Server** Select the EMS server for which you want to view data from this drop-down menu. The selection made here populates this display.
- Name** The name of the EMS Server selected from the EMS Server drop-down menu.
- Server Activity Trends** Specifies settings for the trend graphs.

**Trend Graphs**

Shows metrics for the selected server.

**Connections** -- Traces the total number of client connections.

**Msgs In/Sec** -- Traces the number of inbound messages, per second, from all producers and consumers.

**Msgs Out/Sec** -- Traces the number of outbound messages, per second, from all producers and consumers.

**Pending Msgs** -- Traces the total number of messages currently waiting to be processed.

**Msg Memory %** -- Traces the amount of memory, in percent, used by messages.

**Async Store MB** -- Traces the amount of database space, in megabytes, used by asynchronous data on the server.

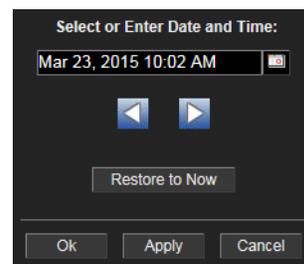
**Disk Read KB** -- Traces the amount of disk data, in kilobytes, read by the server since the server was started.

**Disk Write KB** -- Traces the amount of data, in kilobytes, written to disk by the server since the server was started.

**Log Scale** This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.

**Base at Zero** When this option is checked, zero is set as the Y axis minimum for all graph traces.

**Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



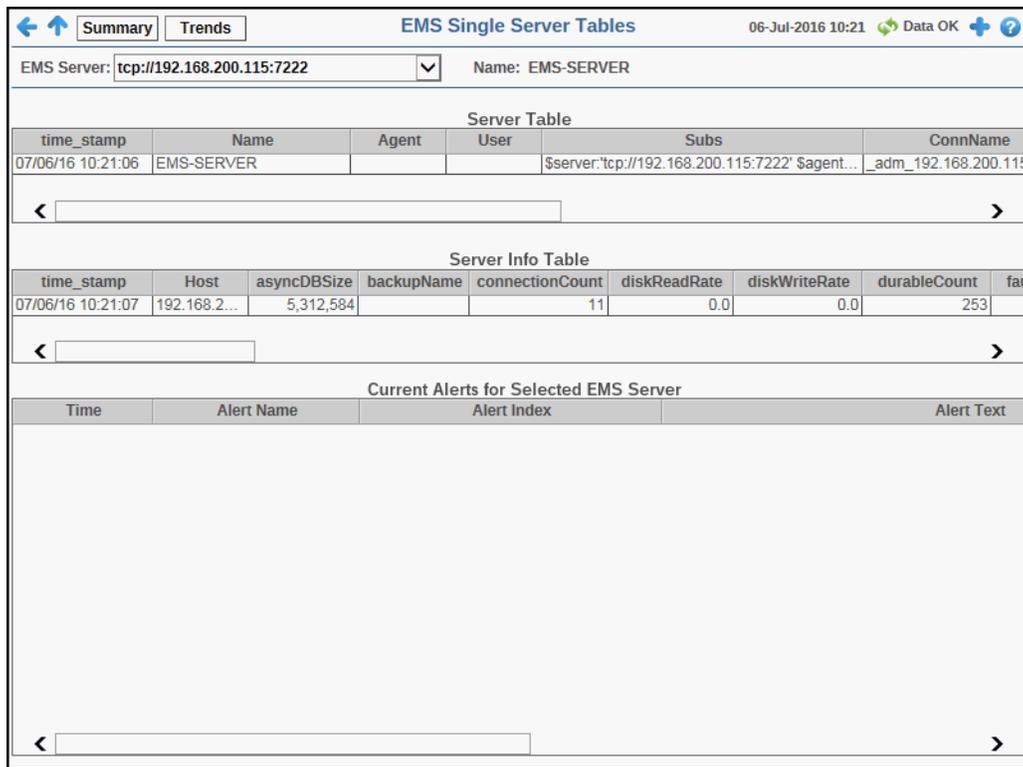
By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

**Single Server Tables**

View all available utilization and performance data for specific servers.



**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** , **Table** open commonly accessed displays.
- 6,047** The number of items currently in the display.
- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

**Fields and Data**

This display includes:

- EMS Server** Select the EMS server for which you want to view data from this drop-down menu. The selection made here populates this display
- Name** The name of the EMS Server selected from the EMS Server drop-down menu.
- Server Table** This table shows information about how the monitor is connected to the server.
  - time\_stamp** The date and time this row of data was last updated.
  - Name** The name of the server.

<b>Agent</b>	If used, the name of the RTView agent connecting to the EMS server.
<b>User</b>	The user name for gaining access to the server.
<b>Password</b>	The password associated with user name for gaining access to the server.
<b>Subs</b>	RTView substitutions used when connecting to this server.
<b>ConnName</b>	The name of the RTView connection to this server.
<b>Active</b>	When checked, indicates that the server is currently running.
<b>FaultTolerantStandbyMode</b>	When checked, indicates that the server is running as a backup server.
<b>FaultTolerantURL</b>	The IP address and port number for the backup server assigned to this server.
<b>BackupName</b>	The name of the backup server assigned as backup to this server.
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > ( <b>Project Name</b> ) > <b>Solution Package Configuration</b> > <b>TIBCO Enterprise Message Service</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.

#### Server Info Table

Select an EMS Server from the EMS Server drop-down menu. This table shows server metrics queried from the server.

<b>time_stamp</b>	The date and time this row of data was last updated.
<b>Host</b>	The name or IP address for the host server.
<b>asyncDBSize</b>	The amount of database space, in bytes, used by asynchronous data on the server.
<b>backupName</b>	The name of the backup server assigned as backup to this server.
<b>connectionCount</b>	The number of currently connected clients.
<b>diskReadRate</b>	The speed at which the server reads disk data.
<b>diskWriteRate</b>	The speed at which the server writes data to disk.
<b>durableCount</b>	The number of currently active durables.
<b>FaultTolerantURL</b>	The IP address and port number, or the hostname and port number, of the fault tolerant standby server assigned to this server.
<b>inboundBytesRate</b>	The rate of inbound messages in bytes per second.
<b>inboundMessageCount</b>	The number of inbound messages received by the server since the server was started.
<b>inboundMessageRate</b>	The rate of inbound messages in number of messages per second.
<b>maxMessageMemory</b>	The maximum amount of memory, in bytes, allocated for use by messages on the server.
<b>messageMemory</b>	The amount of memory, in bytes, currently used by messages on the server.

<b>messageMemoryPct</b>	The amount of memory, in percent, used by messages on the server.
<b>messageMemoryPooled</b>	The currently allocated pool size for messages in bytes.
<b>outboundBytesRate</b>	The rate of outbound messages in bytes per second.
<b>outboundMessageCount</b>	The number of outbound messages sent by the server since the server was started.
<b>outboundMessageRate</b>	The rate of outbound messages in number of messages per second
<b>pendingMessageCount</b>	The number of currently pending messages on the server.
<b>pendingMessageSize</b>	The amount of space, in bytes, pending messages use on the server.
<b>processId</b>	The process ID of the EMS server.
<b>queueCount</b>	The number of message queues.
<b>serverName</b>	The name of the server.
<b>startTime</b>	The date and time that the server was started.
<b>state</b>	The server status: <b>Active</b> -- The server is currently processing requests. <b>Inactive</b> --The server is not currently processing requests. <b>Standby</b> -- The server is functioning as a backup for a primary server.
<b>syncDBSize</b>	The amount of database space, in bytes, used by synchronous data on the server.
<b>topicCount</b>	The number of currently active topics.
<b>upTime</b>	The amount of time, in milliseconds, since the server was started.
<b>versionInfo</b>	The TIBCO EMS software version currently running.
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > ( <b>Project Name</b> ) > <b>Solution Package Configuration</b> > <b>TIBCO Enterprise Message Service</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Current Alerts Table for Selected EMS Server</b>	Select an EMS Server from the EMS Server drop-down menu. This table lists all available data for currently active alerts. Click an alert to view details in the Alert Detail Window.
<b>Time</b>	The time the alert was first activated.
<b>Alert Name</b>	The name of the alert.
<b>Alert Index</b>	The EMS server that activated the alert.
<b>Alert Text</b>	The text that is displayed for the alert.
<b>Package</b>	The RTView package reporting the alert.
<b>Category</b>	The alert category: Server, Queue or Topic.

<b>ID</b>	The unique identifier for this alert instance.
<b>Clr'd</b>	When checked, the alert thresholds are no longer out of bounds and the alert has cleared.
<b>Ack'd</b>	When checked, a user has indicated that they have acknowledged the alert.
<b>Owner</b>	The user who has accepted ownership of this alert.
<b>Source</b>	The source of the alert.

### Alert Detail Window

The screenshot shows the 'Alert Detail' window with the following fields and values:

- Alert Time:** 02/26/15 07:27:45
- ID:** 1000
- Name:** EmsServerMemUsedHigh
- Index:** tcp://SLHOST21:7222
- Owner:** (empty)
- Alert Text:** High Warning Limit exceeded, current value: 9.36 limit: 5.0
- Comments:** (empty text area)
- Severity:** 1
- Acknowledged:**
- Cleared:**

<b>Alert Time</b>	The time the alert was first activated.
<b>ID</b>	The unique identifier for this alert instance.
<b>Name</b>	The name of the alert.
<b>Index</b>	The EMS server which activated the alert.
<b>Owner</b>	The user who has accepted ownership of this alert.
<b>Alert Text</b>	The text that is displayed for the alert.
<b>Comments</b>	User-supplied comments about this alert.
<b>Acknowledged</b>	When checked, a user has indicated that they have acknowledged the alert.
<b>Cleared</b>	When checked, the alert thresholds are no longer out of bounds and the alert has cleared.
<b>Severity</b>	Severity of the alert.

## EMS Topics

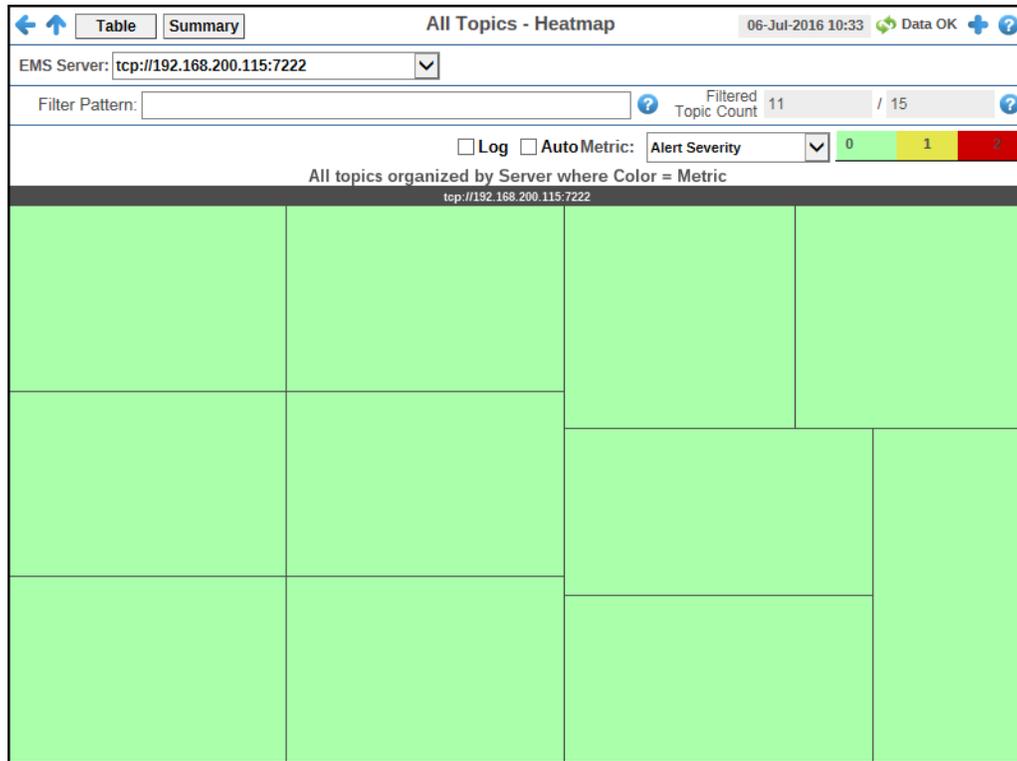
These displays present several views of performance metrics for topics. You can view all topics that are defined on a specific server in the “[All Topics Table](#)” display, or you can view all servers that have a specific topic defined in the “[Single Topic Summary](#)” display. The “[Single Topic By Server](#)” display provides a list of all the servers on which those topics are defined.

- “[All Topics Heatmap](#)”: A heatmap representation of a selected set of metrics from Topics organized by Server that allows you to track performance and utilization metrics and trends for all topics on a single server.
- “[All Topics Table](#)”: Shows performance and utilization metrics and trends for all topics defined on a specified server, including consumer and subscriber count, memory utilization, and message performance metrics.
- “[All Topics Summary](#)”: Shows performance and utilization metrics and trends for all topics defined on a specified server, including consumer and subscriber count, memory utilization, and message performance metrics.
- “[Single Topic Summary](#)”: Shows detailed performance and utilization metrics and trends for a specified topic on a single server, including producer and consumer counts, and message performance metrics.
- “[Single EMS Topic-Clients](#)”: View data for all consumers and producers associated with the selected topic.
- “[Single Topic By Server](#)”: Table shows performance and utilization metrics for all servers that have a specified topic defined, including consumer and subscriber count, and message performance metrics.

### All Topics Heatmap

A heatmap representation of a selected set of metrics from Topics organized by Server that allows you to track performance and utilization metrics and trends for all topics on a single server. View status and alerts of all topics for a server. Use the **Metric** drop-down menu to view to **Alert Severity, Alert Count, Consumers, Receivers, Pending Messages, Inbound Message Rate, Inbound Total Messages, Outbound Message Rate, or Outbound Total Messages**.

The heatmap is organized so that each rectangle represents a Topic on the selected Server. The rectangle color indicates the value of the selected metric in the **Metric** drop down list. You can mouse-over rectangles to view more details about the performance and status of each topic or click on a rectangle to drill-down to the “[Single Topic Summary](#)” display and view metrics for that particular Topic. You can click **Table** on this display to navigate to the “[All Topics Table](#)” display.



#### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** , **Table** open commonly accessed displays.
- 6,047** The number of items currently in the display.

- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

**Note:** Clicking **Table** in the Title Bar takes you to the "All Topics Table" display. Clicking **Summary** in the Title Bar takes you to the "All Topics Summary" display.

### Fields and Data

This display includes:

- EMS Server** The EMS Server selected from this drop-down menu populates all associated Topic data in this display.
- Filter Pattern** Enter a string to show only topics with names that contain the string. For example, if you enter the string Madrid, all topics with Madrid in the topic name are shown in the table. If no entry is made, all topic names are shown. For most use cases, you can enter a portion of the topic name.

**Filtered Topic Count**

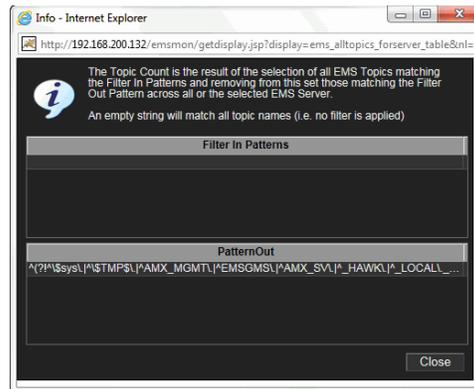
This field is broken into two different values. The first value is the total number of currently active topics on the selected server, which is filtered by the **Filter Pattern** field and by the default value specified in the **\$emsTopicFilterOutPattern** property in the **emsmom/conf/rtvapm.properties** file. The second value is the total number of topics on the selected server. In other words, the filtered number of topics/the total number of topics on the server.

The default value for the **\$emsTopicFilterOutPattern** property is:

```
collector.sl.rtvew.sub=$emsTopicFilterOutPattern:'^(?!\\$sys\\.|^\\$TMP\\$\\.|^AMX_MGMT\\.|^EMSGMS\\.|^AMX_SV\\.|^_HAWK\\.|^_LOCAL\\.|^_HAWK\\.|^TMP\\.EMS)'
```

You can modify the filter value by editing the **\$emsTopicFilterOutPattern** property in the **sample.properties File**, which will override the default value.

Clicking the associated Help button  displays the **Info** dialog, which displays the defined filter in and filter out properties used by the **Filtered Topic Count**.

**Log**

This option enables visualization on a logarithmic scale, and should be used when the range in your data is very broad. For example, if you have data that ranges from the tens to the thousands, then data in the range of tens will be neglected visually if you do not check this option. This option makes data on both extreme ranges visible by using the logarithmic of the values rather than the actual values.

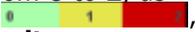
**Auto**

When checked, the values of the selected metric are auto-scaled to its highest defined value. When unchecked, the values of the selected metric display based on the threshold defined for the alert associated with the selected metric. Selecting Auto helps to visualize the range of the values currently present for the selected metric instead of the threshold of the alert that has been associated with the metric. All metrics that have not been associated in the heatmap defaults with alerts use a monochromatic color gradient bar (whites and greens). All metrics that have been associated in the heatmap defaults with alerts use a multi-chromatic color gradient bar (reds, yellows, white, and greens).

**Metric**

Select the metric driving the heatmap display. The default is Alert Severity. Each **Metric** has a color gradient bar that maps values to colors. The heatmap organizes the topics by server, where each rectangle represents a Topic. Mouse-over any rectangle to display the current values of the metrics for the Topic. Click on a rectangle to drill-down to the associated **Single Topic Summary** display for a detailed view of metrics for that particular topic.

**Alert Severity**

The maximum alert level in the item (index) associated with the rectangle. Values range from **0** to **2**, as indicated in the color gradient bar , where **2** is the greatest **Alert Severity**.

**2** -- Metrics that have exceeded their specified **ALARMLEVEL** threshold and have an Alert Severity value of **2** are shown in red. For a given rectangle, this indicates that one or more metrics have exceeded their alarm threshold.

**1** -- Metrics that have exceeded their specified **WARNINGLEVEL** threshold and have an Alert Severity value of **1** are shown in yellow. For a given rectangle, this indicates that one or more metrics have exceeded their warning threshold.

**0** -- Metrics that have not exceeded either specified threshold have an Alert Severity value of **0** and are shown in green. For a given rectangle, this indicates that no metrics have exceeded a specified alert threshold.

**Alert Count**

The total number of alarm and warning alerts in a given item (index) associated with the rectangle.

The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

**Consumers**

The total number of consumers in a given item (index) associated with the rectangle. The color gradient bar  shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of consumers in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

The **Auto** option does not impact this metric.

**Durables**

The total number of active durables in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of durables in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

**Subscribers**

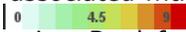
The total number of subscribers in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of subscribers in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

**Pending Msgs**

The total number of pending messages in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of **EmsTopicssPendingMsgsHigh**, which is **3000**. The middle value in the gradient bar indicates the middle value of the range (the default is **1500**).

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

**In Msg /sec**

The number of inbound messages per second in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of **EmsTopicsInMsgRateHigh**, which is **9**. The middle value in the gradient bar indicates the middle value of the range (the default is **4.5**).

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

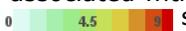
**Note:** This metric comes directly from the **tibjms.admin.DestinationInfo** class from TIBCO.

**In Total Msg**

The total number of inbound messages in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of receivers in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

The **Auto** option does not impact this metric.

**Out Msg/sec**

The number of outbound messages per second in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of **EmsTopicsOutMsgRateHigh**, which is **9**. The middle value in the gradient bar indicates the middle value of the range (the default is **4.5**).

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

**Note:** This metric comes directly from the **tibjms.admin.DestinationInfo** class from TIBCO.

**Out Total Msgs**

The total number of outbound messages in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of receivers in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

The **Auto** option does not impact this metric.

## All Topics Table

Track performance and utilization metrics for all topics on a single server.

Topic Name	URL	In Rate	In Total	Out Rate	Out Total	Pend Msgs	P
adb.custom.jmsrequest	tcp://192.1...	0	0	0	0	0	
adb.salesorder.rr	tcp://192.1...	0	0	0	0	0	
adb.salesorder.sub	tcp://192.1...	0	0	0	0	0	
adb.standard.jmsrequest	tcp://192.1...	0	582,832	0	583,934	2,233,301	1.9
MessageSelector	tcp://192.1...	0	0	0	0	0	
rtv.amx.governance.internal.stats	tcp://192.1...	0	0	0	0	0	
rtv.amx.governance.stats	tcp://192.1...	0	0	0	0	0	
sample	tcp://192.1...	0	0	0	0	0	
topic.sample	tcp://192.1...	0	0	0	0	0	
topic.sample.exported	tcp://192.1...	0	0	0	0	0	
topic.sample.imported	tcp://192.1...	0	0	0	0	0	

**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** dropdown, **Table** button open commonly accessed displays.
- 6,047** The number of items currently in the display.
- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

**Note:** Clicking **Heatmap** in the Title Bar takes you to the "All Topics Heatmap" display. Clicking **Summary** in the Title Bar takes you to the "All Topics Summary" display.

### Fields and Data

This display includes:

- EMS Server** The EMS Server selected from this drop-down menu populates all associated Topic data in this display.

**Filter Pattern** Enter a string to show only topics with names that contain the string. For example, if you enter the string Madrid, all topics with Madrid in the topic name are shown in the table. If no entry is made, all topic names are shown. For most use cases, you can enter a portion of the topic name.

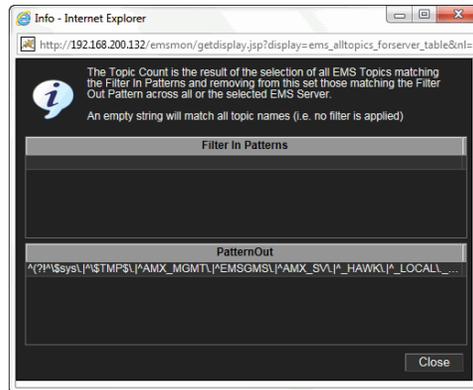
**Filtered Topic Count** This field is broken into two different values. The first value is the total number of currently active topics on the selected server, which is filtered by the **Filter Pattern** field and by the default value specified in the **\$emsTopicFilterOutPattern** property in the **emsmom/conf/rtvapm.properties** file. The second value is the total number of topics on the selected server. In other words, the filtered number of topics/the total number of topics on the server.

The default value for the **\$emsTopicFilterOutPattern** property is:

```
collector.sl.rtvew.sub=$emsTopicFilterOutPattern: '^(?!\\$sys\\.|^\\$TMP\\$\\.|^AMX_MGMT\\.|^EMSGMS\\.|^AMX_SV\\.|^_HAWK\\.|^_LOCAL\\. _HAWK\\.|^TMP\\.EMS)'
```

You can modify the filter value by editing the **\$emsTopicFilterOutPattern** property in the **sample.properties File**, which will override the default value.

Clicking the associated Help button  displays the **Info** dialog, which displays the defined filter in and filter out properties used by the **Filtered Topic Count**.



**Table** This table describes all topics on the selected server. Click a row to view metrics for a single topic in the **Single Topic Summary** display.

<b>Topic Name</b>	The name of the topic.
<b>URL</b>	The IP address and port number for the server.
<b>In Rate</b>	The number of inbound messages for the topic, per second. <b>Note:</b> This metric comes directly from the <b>tibjms.admin.DestinationInfo</b> class from TIBCO.
<b>In Total</b>	The total number of inbound messages for the topic.
<b>Out Rate</b>	The number of outbound messages for the topic, per second. <b>Note:</b> This metric comes directly from the <b>tibjms.admin.DestinationInfo</b> class from TIBCO.
<b>Out Total</b>	The total number of outbound messages for the topic.
<b>Pend Msgs</b>	The number of currently pending messages for the topic.
<b>Pend Size</b>	The amount of space, in bytes, used by pending messages for the topic.
<b>activeDurableCount</b>	The number of currently active durables or the topic.

<b>consumerCount</b>	The number of consumers for the topic.
<b>durableCount</b>	The number of durables for the topic.
<b>failSafe</b>	When checked, the message is marked as failsafe delivery.
<b>fcMaxBytes</b>	The maximum number of bytes allocated for use by flow control.
<b>global</b>	When checked, the message is global and is routed to other servers.
<b>inboundByteRate</b>	The amount of inbound messages for the topic, in bytes per second.
<b>inboundTotalBytes</b>	The total amount of inbound messages for the topic, in bytes, since the server started.
<b>maxBytes</b>	The maximum size, in bytes, that the topic can store for delivery to each durable or non-durable online subscriber on that topic.
<b>maxMsgs</b>	The maximum number of messages before the server indicates an error and overflow policies are activated.
<b>outboundByteRate</b>	The amount of outbound messages for the topic, in bytes per second.
<b>outboundTotalBytes</b>	The total amount of outbound messages for the topic, in bytes.
<b>overflowPolicy</b>	Indicates whether an overflow policy is set for the topic: <b>0</b> = No policy is set. <b>1</b> = A policy is set.
<b>secure</b>	When checked, the topic is designated as secure and enforces permission policies.
<b>static</b>	When checked, the topic has a static destination.
<b>subscriberCount</b>	The number of subscribers for the topic.
<b>description</b>	Descriptive text to help the administrator identify this resource.
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > <b>(Project Name)</b> > <b>Solution Package Configuration</b> > <b>TIBCO Enterprise Message Service</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>time_stamp</b>	The date and time this row of data was last updated.
<b>DeltainboundTotalMessages</b>	Displays the change (delta) in inboundTotalMessages from the previous cache refresh to the current cache refresh.
<b>DeltainboundTotalBytes</b>	Displays the change (delta) in inboundTotalBytes from the previous cache refresh to the current cache refresh.
<b>DeltaoutboundTotalMessages</b>	Displays the change (delta) in outboundTotalMessages from the previous cache refresh to the current cache refresh.

**DeltaoutboundTotalBytes**

Displays the change (delta) in outboundTotalBytes from the previous cache refresh to the current cache refresh.

**prefetch**

Lists the maximum number of messages consumers can fetch.

**expiryOverride**

If set to a non-zero value for a destination and the server delivers a message to the destination, the server replaces the producer's expiration value with this value.

**store**

Provides the store for this destination where persistent messages are stored.

**URLTopic**

The topic's URL.

### All Topics Summary

Track performance and utilization metrics and trends for all topics on a single server.



**Title Bar** (possible features are):

-   Open the previous and upper display.
-  Open an instance of this display in a new window.
-  Open the online help page for this display.
- ,  open commonly accessed displays.
- The number of items currently in the display.

 **Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

 Open the **Alert Views - RTView Alerts Table** display.

---

**Note:** Clicking **Heatmap** in the Title Bar takes you to the "All Topics Heatmap" display. Clicking **Table** in the Title Bar takes you to the "All Topics Table" display.

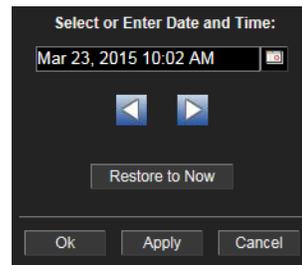
---

### Fields and Data

This display includes:

<b>EMS Server</b>	The EMS Server selected from this drop-down menu populates all associated Topic data in this display.
<b>Name</b>	The name of the server selected in the <b>EMS Server</b> drop down list.
<b>Totals for Server</b>	Shows metrics for all topics on the selected server.
<b>Messages In</b>	<p><b>Msgs/sec</b> -- The number of inbound messages for all topics on the server, per second.</p> <p><b>Total</b> -- The total number of inbound messages for all topics on the server since the server was started.</p> <p><b>Bytes/sec</b> -- The size of inbound messages, in bytes per second, for all topics on the server.</p> <p><b>Total</b> -- The total size of inbound messages, in kilobytes, for all topics on the server since the server was started.</p>
<b>Messages Out</b>	<p><b>Msgs/sec</b> -- The number of outbound messages for all topics on the server, per second.</p> <p><b>Total</b> -- The total number of outbound messages for all topics on the server since the server was started.</p> <p><b>Bytes/sec</b> -- The size of outbound messages, in bytes per second, for all topics on the server.</p> <p><b>Total</b> -- The total size of outbound messages for all topics on the server, in kilobytes, since the server was started.</p>
<b>Pending Messages</b>	<p><b>Current</b> -- The total number of messages for all topics on the server currently waiting to be processed.</p> <p><b>Bytes</b> -- The total size of messages, in bytes, for all topics on the server currently waiting to be processed.</p>
<b>Trend Graphs</b>	Shows metrics for all topics on the selected server. <ul style="list-style-type: none"> <li><b>Pend Msgs</b> -- Traces the total number of messages for all topics on the server currently waiting to be processed.</li> <li><b>In Msgs / sec</b> -- Traces the number of inbound messages for all topics, per second.</li> <li><b>Out Msgs / sec</b> -- Traces the number of outbound messages for all topics, per second.</li> </ul>

- Log Scale** This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.
- Base at Zero** When this option is checked, zero is set as the Y axis minimum for all graph traces.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



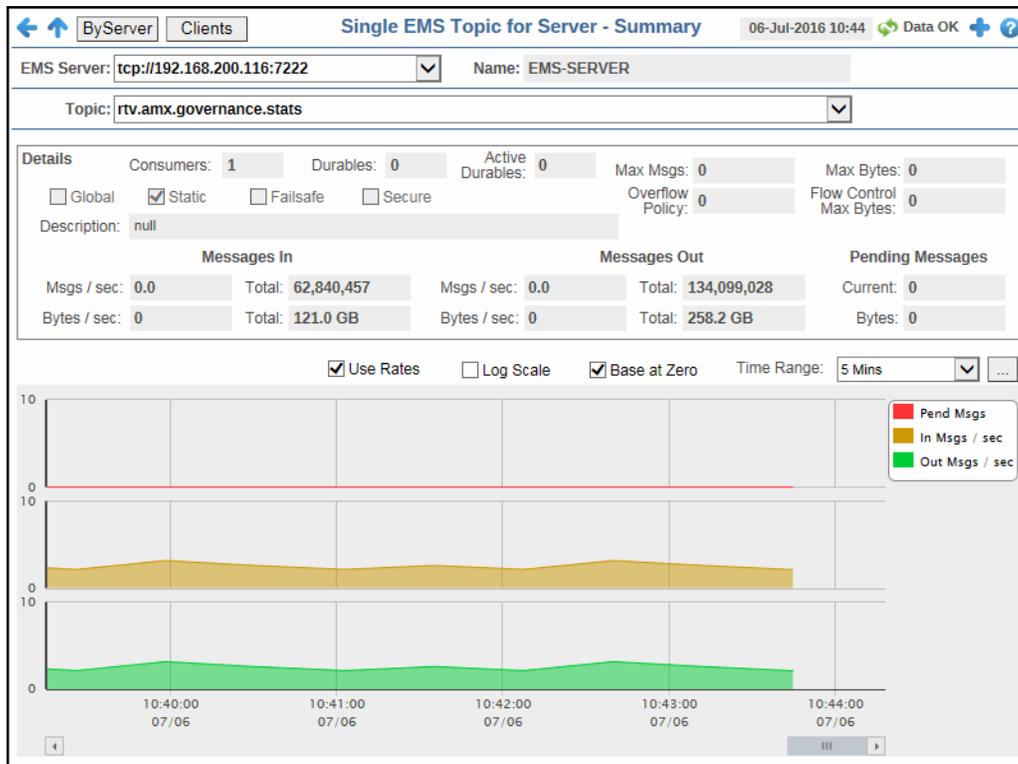
By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Single Topic Summary

Track performance and utilization metrics for a single topic on a single server.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** **Table** open commonly accessed displays.
- 6,047** The number of items currently in the display.

- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

**Note:** Clicking **Clients** in the Title Bar takes you to the "Single EMS Topic-Clients" display for the selected topic.

### Fields and Data

This display includes:

- EMS Server** The EMS Server selected from this drop-down menu populates the Topics drop-down menu with the Topics belonging to this EMS Server.
- Name** The name of the EMS server selected from the EMS Server drop-down menu.

**Topic** Select a Topic from the drop-down menu to view details for the selected Topic.

**Browse** Click to browse the contents of the selected topic in a separate window. The topic browser table displays up to 100,000 rows of messages.

Topic	MsgID	Browse#	Status
Alca Chan	5272001	230	BROKEN
Alca Chan	5272001	229	WORKING
Alca Chan	5272001	228	WORKING
Alca Chan	5272001	227	WORKING
Alca Chan	5272001	226	WORKING
Alca Chan	5272001	225	WORKING
Alca Chan	5272001	224	WORKING
Alca Chan	5272001	223	WORKING
Alca Chan	5272001	222	WORKING
Alca Chan	5272001	221	WORKING
Alca Chan	5272001	220	WORKING
Alca Chan	5272001	219	WORKING
Alca Chan	5272001	218	WORKING

By default, this button is disabled due to the fact that use of this option could significantly impact performance. To enable it, add the following substitution to the properties file with which you execute the Display Server and/or Viewer:

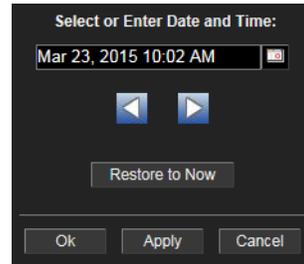
**sl.rtview.sub=\$emsDestBrowseButtonVisFlag:1**

**Details** Shows metrics for the topic selected from the Topic drop-down menu.

- Consumers** The current number of consumers for the topic.
- Durables** The number of durable subscribers (active and inactive) to the topic.
- Active Durables** The number of active durable subscribers to the topic.
- Max Msgs** The maximum number of messages allocated for the topic.
- Max Bytes** The maximum of memory, in bytes, allocated for use by the topic.
- Global** When checked, the message is global and is routed to other servers.
- Static** When checked, the topic has a static destination.
- Failsafe** When checked, the message is marked as failsafe delivery.
- Secure** When checked, the topic is designated as secure and enforces permission policies.
- Overflow Policy** Indicates whether an overflow policy is set for the topic:  
**0** = No policy is set.  
**1** = A policy is set.
- Flow Control Max Bytes** The maximum amount of memory, in bytes, allocated for flow control use by the topic.
- Description** Description of the Topic.
- Messages In**
  - Msgs/sec** The number of inbound messages, per second, for the selected topic.
  - Total** The total number of inbound messages for the selected topic since the server was started.
  - Bytes/sec** The size of inbound messages, in bytes per second, for the selected topic.
  - Total** The total size of inbound messages, in bytes, for the selected topic since the server was started.
- Messages Out**
  - Msgs/sec** The number of outbound messages, per second, for the selected topic.
  - Total** The total number of outbound messages for the selected topic since the server was started.
  - Bytes/sec** The size of outbound messages, in bytes per second, for the selected topic.

	<b>Total</b>	The total size of outbound messages, in bytes, for the selected topic since the server was started.
<b>Pending Messages</b>	<b>Current</b>	The number of messages for the selected topic currently waiting to be processed.
	<b>Bytes</b>	The size of the messages for the selected topic, in bytes, currently waiting to be processed.
<b>Trend Graphs</b>		Shows message data for the selected topic.
		<b>Pend Msgs</b> -- Traces the number of messages currently waiting to be processed.
		<b>In Msgs / sec</b> -- Traces the number of inbound messages, per second. This trend graph only displays when <b>Use Rates</b> is selected.
		<b>Out Msgs / sec</b> -- Traces the number of outbound messages, per second. This trend graph only displays when <b>Use Rates</b> is selected.
		<b>Delta In Msgs</b> -- Traces the change in total inbound messages since the last update. This trend graph only displays when <b>Use Rates</b> is not selected.
		<b>Delta Out Msgs</b> -- Traces the change in total inbound messages since the last update. This trend graph only displays when <b>Use Rates</b> is not selected.
		<b>Use Rates</b> When this check box is selected, the inbound and outbound message rates ( <b>In Msgs/sec</b> and <b>Out Msgs/sec</b> ) display in the trend graph. When this check box is not selected, the delta inbound and outbound messages ( <b>Delta In Msgs</b> and <b>Delta Out Msgs</b> ) display in the trend graph.
	<b>Log Scale</b> This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.	

- Base at Zero** When this option is checked, zero is set as the Y axis minimum for all graph traces.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Single EMS Topic-Clients

View data for all consumers and producers associated with the selected topic.

← ↑ ByServer Summary **Single EMS Topic - Clients** 06-Jul-2016 10:51 Data OK

EMS Server: tcp://192.168.200.116:7222 Name: EMS-SERVER

Topic: rtv.amx.governance.stats  Show Active Only

Producers Count: 0

ID	clientID	Msgs / sec	Msgs Total	Bytes / sec	Total Bytes	userName	hos

Consumers Count: 1

ID	clientID	Msgs/sec	Msgs Total	Bytes/sec	Total Bytes	userName	hos
5614606		0.0	1,739,409	0.0	3,626,981,...	admin	SLHOST21

### Title Bar (possible features are):

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu ▾, Table open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

**Note:** Clicking **Summary** in the Title Bar takes you to the "Single Topic Summary" display. Clicking **ByServer** in the Title Bar takes you to the "Single Topic By Server" display.

### Fields and Data

This display includes:

- EMS Server** The EMS Server selected from this drop-down menu populates the Topics drop-down menu with the Topics belonging to this EMS Server.
- Name** The name of the EMS Server selected from the EMS Server drop-down menu.

<b>Topic</b>	Select a Topic from the drop-down menu to view details for the selected Topic.
<b>Show Active Only</b>	Select this check box to view only the active producers and consumers for the selected Server/ Topic combination.
<b>Producers</b>	Shows data for all producers for the selected topic.
<b>ID</b>	A unique string identifier assigned to each producer.
<b>clientID</b>	A unique string identifier assigned to each client.
<b>Msgs / sec</b>	The number of messages, per second, emitted by the producer.
<b>Msgs Total</b>	The total number of messages emitted by the producer since the server was started.
<b>Bytes / sec</b>	The size of messages, in bytes per second, emitted by the producer.
<b>Total Bytes</b>	The total size of messages, in bytes, emitted by the producer since the server was started.
<b>userName</b>	The user name.
<b>host</b>	The name of the host.
<b>sessionID</b>	A unique string identifier assigned to each session.
<b>connection ID</b>	A unique string identifier assigned to each connection.
<b>createTime</b>	The amount of time, in milliseconds, since the producer was created.
<b>time_stamp</b>	The date and time this row of data was last updated.
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > <b>(Project Name) &gt; Solution Package Configuration &gt; TIBCO Enterprise Message Service &gt; DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Consumers</b>	Shows data for all consumers of messages for the selected topic.
<b>ID</b>	A unique string identifier assigned to each consumer.
<b>clientID</b>	A unique string identifier assigned to each client.
<b>Msgs / sec</b>	The number of messages, per second, processed by the consumer.
<b>Msgs Total</b>	The total number of messages processed by the consumer.
<b>Bytes / sec</b>	The size of messages, in bytes per second, processed by the consumer.
<b>Total Bytes</b>	The total size of messages, in bytes, processed by the consumer since the server was started.
<b>userName</b>	The user name.
<b>host</b>	The name of the host machine.
<b>Msgs Sent</b>	The number of messages sent to the consumer that were not yet acknowledged by the consumer's session. The <b>sl.rtvew.jmsadm.queryCIDetails</b> property must be set to <b>true</b> in your <b>sample.properties</b> file to see this column.

<b>Size Msg Sent</b>	<p>The combined size of messages sent to the consumer that were not yet acknowledged by the consumer's session.</p> <p>The <b>sl.rtvview.jmsadm.queryCIDetails</b> property must be set to <b>true</b> in your <b>sample.properties</b> file to see this column.</p>
<b>Ack Msgs</b>	<p>The total number of messages that have been sent to the consumer and have been acknowledged by the consumer's session.</p> <p>The <b>sl.rtvview.jmsadm.queryCIDetails</b> property must be set to <b>true</b> in your <b>sample.properties</b> file to see this column.</p>
<b>Sent Msgs</b>	<p>The total number of messages sent to the consumer since the consumer was created.</p> <p>The <b>sl.rtvview.jmsadm.queryCIDetails</b> property must be set to <b>true</b> in your <b>sample.properties</b> file to see this column.</p>
<b>Elap. Since Last Ack</b>	<p>The amount of time (in milliseconds) that has elapsed since the last time a message sent to the consumer was acknowledged by the consumer's session.</p> <p>The <b>sl.rtvview.jmsadm.queryCIDetails</b> property must be set to <b>true</b> in your <b>sample.properties</b> file to see this column.</p>
<b>Elap. Since Last Sent</b>	<p>The amount of time (in milliseconds) that has elapsed since the last time the server sent a message to the consumer.</p> <p>The <b>sl.rtvview.jmsadm.queryCIDetails</b> property must be set to <b>true</b> in your <b>sample.properties</b> file to see this column.</p>
<b>destination Prefetch</b>	<p>The actual destination prefetch value used by the server at runtime.</p> <p>The <b>sl.rtvview.jmsadm.queryCIDetails</b> property must be set to <b>true</b> in your <b>sample.properties</b> file to see this column.</p>
<b>prefetch Delivered Count</b>	<p>The number of prefetch messages delivered to the consumer by the server. For consumers receiving messages on any destination with positive prefetch value, this value is never more than the prefetch value of the destination. This value cannot be used to identify the status of the consumer, but it can be used in conjunction with other consumer information values to identify consumers who stopped receiving messages due to application-specific problems.</p> <p>The <b>sl.rtvview.jmsadm.queryCIDetails</b> property must be set to <b>true</b> in your <b>sample.properties</b> file to see this column.</p>
<b> durable Name</b>	<p>The name of the durable.</p>
<b>routeName</b>	<p>The queue owner server name if the consumer's destination is a routed queue.</p> <p>The <b>sl.rtvview.jmsadm.queryCIDetails</b> property must be set to <b>true</b> in your <b>sample.properties</b> file to see this column.</p>
<b>isActive</b>	<p>When checked, the consumer is active and can receive messages from the server.</p> <p>The <b>sl.rtvview.jmsadm.queryCIDetails</b> property must be set to <b>true</b> in your <b>sample.properties</b> file to see this column.</p>
<b>isSystem</b>	<p>This check box is checked if the consumer was automatically created by the system.</p> <p>The <b>sl.rtvview.jmsadm.queryCIDetails</b> property must be set to <b>true</b> in your <b>sample.properties</b> file to see this column.</p>
<b>sessionAck Mode</b>	<p>Lists the consumer's session acknowledge mode as a constant defined in <b>TibjmsAdmin</b>.</p> <p>The <b>sl.rtvview.jmsadm.queryCIDetails</b> property must be set to <b>true</b> in your <b>sample.properties</b> file to see this column.</p>
<b>session ID</b>	<p>A unique string identifier assigned to each session.</p>

<b>connection ID</b>	A unique string identifier assigned to each connection.
<b>createTime</b>	The amount of time, in milliseconds, since the consumer was created.
<b>time_stamp</b>	The date and time this row of data was last updated.
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > <b>(Project Name) &gt; Solution Package Configuration &gt; TIBCO Enterprise Message Service &gt; DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.

## Single Topic By Server

Track performance and utilization metrics of a single topic across all servers that have the topic defined on it. Compare topic activity among servers.

Topic: adb.standard.jmsrequest

URL	Act. Durables	Consumers	Durables	failsafe	fcMaxBytes	global	In B
tcp://192.168.200.115:7222	1	8	8	<input type="checkbox"/>	0	<input type="checkbox"/>	

**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- The number of items currently in the display.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

Open the **Alert Views - RTView Alerts Table** display.

---

**Note:** Clicking **Clients** in the Title Bar takes you to the "Single EMS Topic-Clients" display for the selected topic. Clicking **Summary** in the Title Bar takes you to the "Single Topic Summary" display.

---

### Fields and Data

This display includes:

<b>Topic</b>	The Topic selected from this drop-down menu populates this display.
<b>Table</b>	Shows details about the selected Topic for each server that has the Topic defined. Select a server from the list to view details in the "Single Topic Summary" display.
<b>URL</b>	The IP address and port number for the server.
<b>Act. Durables</b>	The number of currently active durables.
<b>Consumers</b>	The current number of consumers.
<b>Durables</b>	The number of active and inactive durables.
<b>failsafe</b>	When checked, the message is marked as failsafe delivery.
<b>fcMaxBytes</b>	The maximum number of bytes allocated for use by flow control.
<b>global</b>	When checked, the message is global and is routed to other servers.
<b>In Byte Rate</b>	The amount of inbound messages for the topic, in bytes per second.
<b>In Msgs Rate</b>	The amount of inbound messages for the topic, in number of messages per second.
<b>In Total Bytes</b>	The total number of inbound bytes for the topic.
<b>In Total Msgs</b>	The total number of inbound messages for the topic.
<b>maxBytes</b>	The maximum size, in bytes, that the topic can store for delivery to each durable or non-durable online subscriber on the topic.
<b>maxMsgs</b>	The maximum number of messages allocated for use by the topic.
<b>Out Byte Rate</b>	The amount of outbound messages (in bytes) per second.
<b>Out Msg Rate</b>	The number of outbound messages per second.
<b>Out Total Bytes</b>	The total amount of outbound messages for the topic, in bytes, since the server was started.

<b>Out Total Msgs</b>	The total number of outbound messages for the topic since the server was started.
<b>overflowPolicy</b>	Policy Indicates whether an overflow policy is set for the topic: <b>0</b> = No policy is set. <b>1</b> = A policy is set.
<b>Pending Msgs</b>	The number of currently pending messages for the topic.
<b>Pending Msgs Size</b>	The amount of space, in bytes, pending messages use for the topic.
<b>secure</b>	When checked, the topic is designated as secure and enforces permission policies.
<b>static</b>	When checked, the topic has a static destination.
<b>Subscribers</b>	The number of subscribers for the topic.
<b>time_stamp</b>	The date and time this row of data was last updated.
<b>description</b>	Descriptive text to help the administrator identify this resource.

## EMS Queues

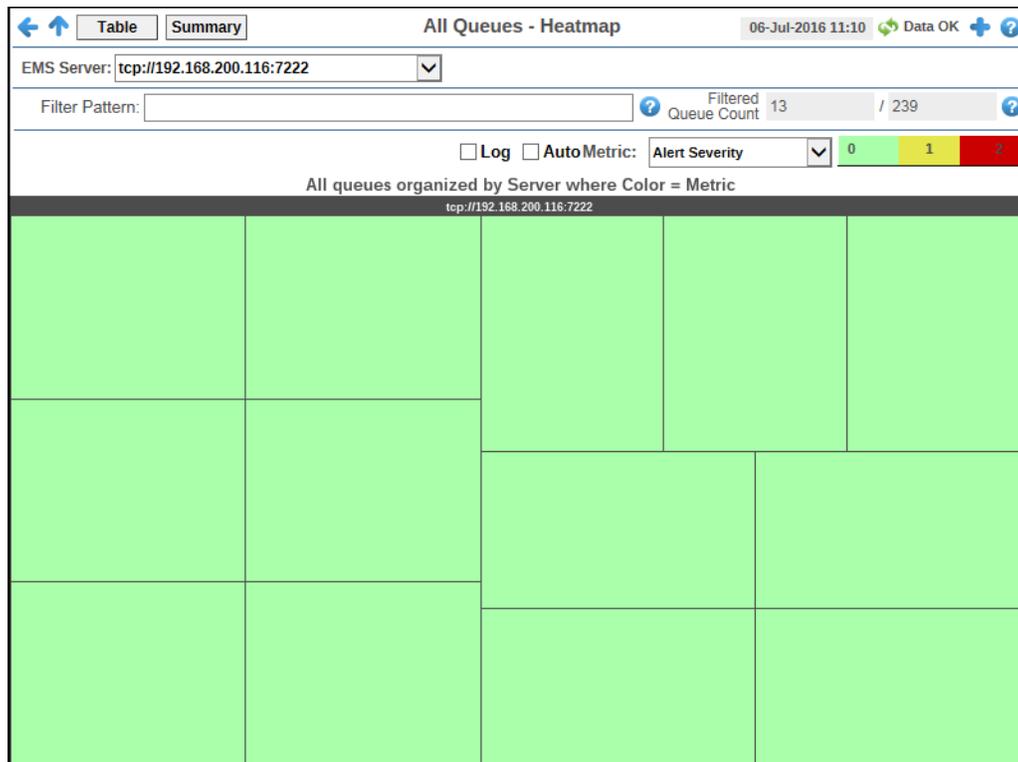
These displays present several views of performance metrics for queues. You can view all queues that are defined on a specific server in the ["All Queues Heatmap"](#) display, or you can view all servers that have a specific queue defined in the ["Single Queue Summary"](#) display. The ["Single EMS Queue-Clients"](#) display provides a list of all the servers on which those queues are defined.

- ["All Queues Heatmap"](#): A heatmap representation of a selected set of metrics that shows performance and utilization metrics and trends for all queues defined on a specified server, including message performance metrics.
- ["All Queues Table"](#): Shows performance and utilization metrics for all queues defined on a specified server.
- ["All Queues Summary"](#): Shows performance and utilization metrics and trends for all queues defined on a specified server, including message performance metrics.
- ["Single Queue Summary"](#): Shows detailed performance and utilization metrics and trends for a specified queue on a single server, including producer and consumer counts, and message performance metrics.
- ["Single EMS Queue-Clients"](#): View data for all consumers and producers associated with the selected queue.
- ["Single Queue By Server"](#): Table shows performance and utilization metrics for all servers that have a specified queue defined, including consumer and receiver count, and message performance metrics.

## All Queues Heatmap

A heatmap representation of the ["All Queues Table"](#) display that allows you to track performance and utilization metrics and trends for all queues on a single server. View status and alerts of all queues for a server. Use the **Metric** drop-down menu to view to **Alert Severity, Alert Count, Consumers, Receivers, Pending Messages, Inbound Message Rate, Inbound Total Messages, Outbound Message Rate, or Outbound Total Messages**.

The heatmap is organized so that each rectangle represents a queue on the selected server. The rectangle color indicates the most critical alert state. Click on a node to drill-down to the ["Single Queue Summary"](#) display and view metrics for a particular queue. Toggle between the commonly accessed **Table** (link to the ["All Queues Table"](#) display) and **Heatmap** displays. Mouse-over rectangles to view more details about the performance and status of each queue.



#### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** , **Table** open commonly accessed displays.
- 6,047** The number of items currently in the display.

- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

**Note:** Clicking **Table** in the Title Bar takes you to the ["All Queues Table"](#) display. Clicking **Summary** in the Title Bar takes you to the ["All Queues Summary"](#) display.

#### Fields and Data

This display includes:

- EMS Server** The EMS Server selected from this drop-down menu populates all the associated Queue data in this display.

**Filter Pattern** Enter a string to show only queues with names that contain the string. For example, if you enter the string Madrid, all queues with Madrid in the queue name are shown in the table. If no entry is made, all queue names are shown. For most use cases, you can enter a portion of the queue name.

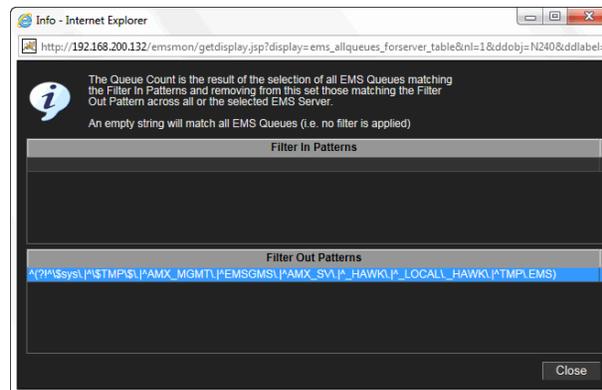
**Filtered Queue Count** This field is broken into two different values. The first value is the total number of currently active queues on the selected server, which is filtered by the **Filter Pattern** field and by the default value specified in the **\$emsQueueFilterOutPattern** property in the **emsmmon/conf/rtvapm.properties** file. The second value is the total number of queues on the selected server. In other words, the filtered number of queues/the total number of queues on the server.

The default value for the **\$emsQueueFilterOutPattern** property is:

```
collector.sl.rtvew.sub=$emsQueueFilterOutPattern: '^(?!\\$sys\\.|^\\$TMP\\$\\.|^AMX_MGMT\\.|^EMSGMS\\.|^AMX_SV\\.|^_HAWK\\.|^_LOCAL\\. _HAWK\\.|^TMP\\.EMS)'
```

You can modify the filter value by editing the **\$emsQueueFilterOutPattern** property in the **"sample.properties File"**, which will override the default value.

Clicking the associated Help button  displays the **Info** dialog, which displays the defined filter in and filter out properties used by the **Filtered Queue Count**.



**Log** This option enables visualization on a logarithmic scale, and should be used when the range in your data is very broad. For example, if you have data that ranges from the tens to the thousands, then data in the range of tens will be neglected visually if you do not check this option. This option makes data on both extreme ranges visible by using the logarithmic of the values rather than the actual values.

**Auto** When checked, the values of the selected metric are auto-scaled to its highest defined value. When unchecked, the values of the selected metric display based on the threshold defined for the alert associated with the selected metric. Selecting Auto helps to visualize the range of the values currently present for the selected metric instead of the threshold of the alert that has been associated with the metric. All metrics that have not been associated in the heatmap defaults with alerts use a monochromatic color gradient bar (whites and greens). All metrics that have been associated in the heatmap defaults with alerts use a multi-chromatic color gradient bar (reds, yellows, white, and greens).

**Metric** Select the metric driving the heatmap display. The default is **Alert Severity**. Each Metric has a color gradient bar that maps values to colors. The heatmap organizes the topics by server, where each rectangle represents a Queue. Mouse-over any rectangle to display the current values of the metrics for the Queue. Click on a rectangle to drill-down to the associated **"Single Queue Summary"** display for a detailed view of metrics for that particular queue.

**Alert Severity**

The maximum alert level in the item (index) associated with the rectangle. Values range from **0** to **2**, as indicated in the color gradient bar , where **2** is the greatest **Alert Severity**.

-- Metrics that have exceeded their specified **ALARMLEVEL** threshold and have an Alert Severity value of **2** are shown in red. For a given rectangle, this indicates that one or more metrics have exceeded their alarm threshold.

**1** -- Metrics that have exceeded their specified **WARNINGLEVEL** threshold and have an Alert Severity value of **1** are shown in yellow. For a given rectangle, this indicates that one or more metrics have exceeded their warning threshold.

-- Metrics that have not exceeded either specified threshold have an Alert Severity value of **0** and are shown in green. For a given rectangle, this indicates that no metrics have exceeded a specified alert threshold.

**Alert Count**

The total number of alarm and warning alerts in a given item (index) associated with the rectangle.

The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

**Consumers**

The total number of consumers in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of receivers in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

The **Auto** option does not impact this metric.

**Receivers**

The total number of receivers in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of receivers in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

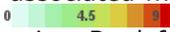
The **Auto** option does not impact this metric.

**Pending Msgs**

The total number of pending messages in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of **EmsQueuesPendingMsgsHigh**, which is **3000**. The middle value in the gradient bar indicates the middle value of the range (the default is **1500**).

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

**In Msgs /sec**

The number of inbound messages per second in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of **EmsQueuesInMsgRateHigh**, which is **9**. The middle value in the gradient bar indicates the middle value of the range (the default is **4.5**).

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

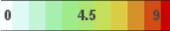
**Note:** This metric comes directly from the **tibjms.admin.DestinationInfo** class from TIBCO.

**In Total Msg**

The total number of inbound messages in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of receivers in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

The **Auto** option does not impact this metric.

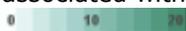
**Out Msgs/sec**

The number of outbound messages per second in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of **EmsQueuesOutMsgRateHigh**, which is **9**. The middle value in the gradient bar indicates the middle value of the range (the default is **4.5**).

When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

**Note:** This metric comes directly from the **tibjms.admin.DestinationInfo** class from TIBCO.

**Out Total Msgs**

The total number of outbound messages in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of receivers in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

The **Auto** option does not impact this metric.

## All Queues Table

Track performance and utilization metrics for all queues on a single server.

Queue Name	URL	In Rate	In Total	Out Rate	Out Total	Pend Msgs
amx.governance.internal.stats	tcp://192.1...	0	206,708,124	0	206,708,124	103
amx.governance.stats	tcp://192.1...	0	62,845,363	0	62,845,367	0
cl_logservice_queue	tcp://192.1...	0	0	0	0	0
cl_payload_queue	tcp://192.1...	0	0	0	0	0
com.tibco.amf.admin.deploymentServerQueue.inst...	tcp://192.1...	0	4	0	4	0
com.tibco.amf.admin.deploymentServerQueue.inst...	tcp://192.1...	0	9	0	9	0
com.tibco.amf.admin.deploymentServerQueue.SL...	tcp://192.1...	0	0	0	0	0
com.tibco.amf.admin.deploymentServerQueue.SL...	tcp://192.1...	0	11	0	11	0
com.tibco.amf.admin.deploymentServerQueue.Sys...	tcp://192.1...	0	0	0	0	0
com.tibco.amf.admin.deploymentServerQueue.Sys...	tcp://192.1...	0	8	0	8	0
com.tibco.amf.admin.deploymentServerQueue.Sys...	tcp://192.1...	0	0	0	0	0
queue.sample	tcp://192.1...	0	0	0	0	0
sample	tcp://192.1...	0	0	0	0	0

### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** , **Table** open commonly accessed displays.
- 6,047** The number of items currently in the display.

- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

**Note:** Clicking **Heatmap** in the Title Bar takes you to the "All Queues Heatmap" display. Clicking **Summary** in the Title Bar takes you to the "All Queues Summary" display.

### Fields and Data

This display includes:

- EMS Server** The EMS Server selected from this drop-down menu populates all associated Queue data in this display.

**Filter Pattern** Enter a string to show only queues with names that contain the string. For example, if you enter the string Madrid, all queues with Madrid in the queue name are shown in the table. If no entry is made, all queue names are shown. For most use cases, you can enter a portion of the queue name.

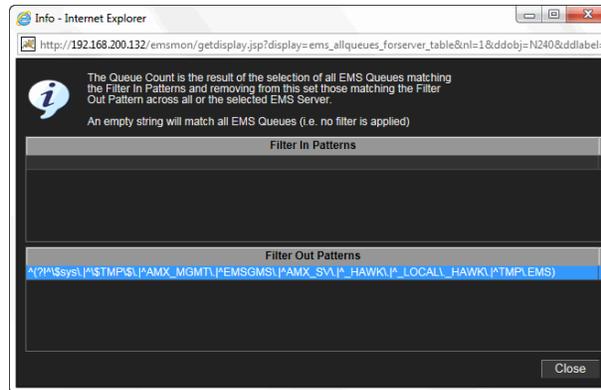
**Filtered Queue Count** This field is broken into two different values. The first value is the total number of currently active queues on the selected server, which is filtered by the **Filter Pattern** field and by the default value specified in the **\$emsQueueFilterOutPattern** property in the **emsmon/conf/rtvapm.properties** file. The second value is the total number of queues on the selected server. In other words, the filtered number of queues/the total number of queues on the server.

The default value for the **\$emsQueueFilterOutPattern** property is:

```
collector.sl.rtvview.sub=$emsQueueFilterOutPattern: '^(?!^\\$sys\\.|^\\$TMP
P\\.|^AMX_MGMT\\.|^EMSGMS\\.|^AMX_SV\\.|^_HAWK\\.|^_LOCAL\\.|^_HAWK\\.|^
^TMP\\.EMS)'
```

You can modify the filter value by editing the **\$emsQueueFilterOutPattern** property in the **sample.properties File**, which will override the default value.

Clicking the associated Help button  displays the **Info** dialog, which displays the defined filter in and filter out properties used by the **Filtered Queue Count**.



**Table** This table describes all queues on the selected server. Click a row to view metrics for a single queue in the **Single Queue Summary** display.

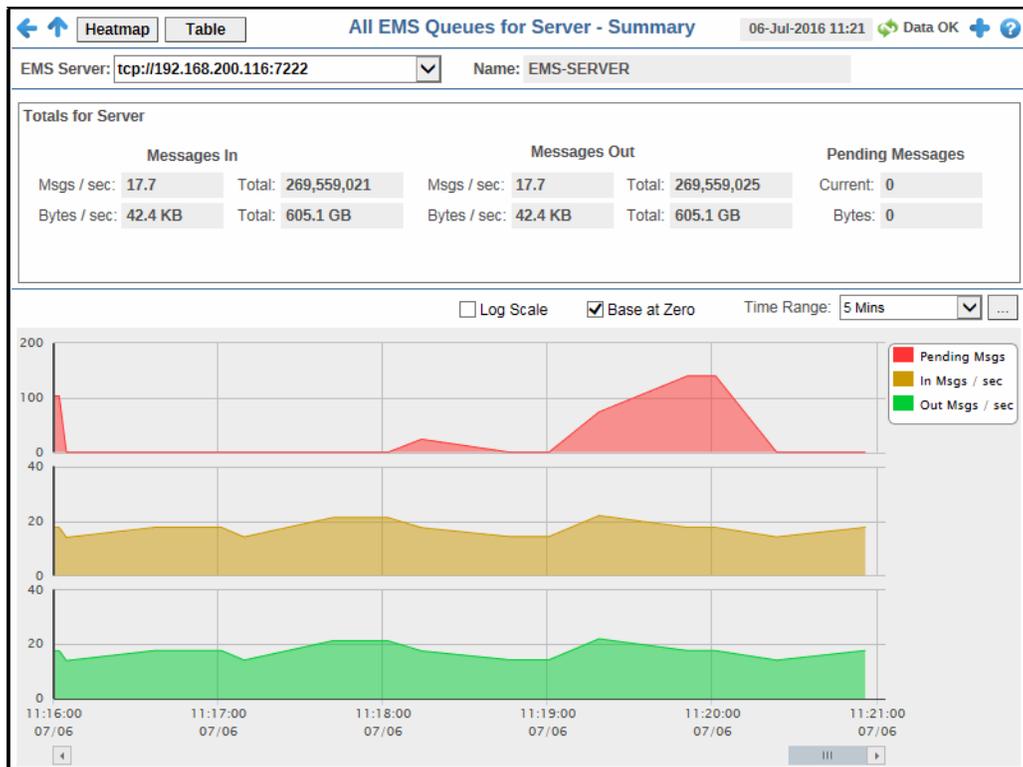
<b>Queue Name</b>	The name of the queue.
<b>URL</b>	The IP address and port number for the server.
<b>In Rate</b>	The number of inbound messages for the queue, per second. <b>Note:</b> This metric comes directly from the <b>tibjms.admin.DestinationInfo</b> class from TIBCO.
<b>In Total</b>	The total number of inbound messages for the queue.
<b>Out Rate</b>	The number of outbound messages for the queue, per second. <b>Note:</b> This metric comes directly from the <b>tibjms.admin.DestinationInfo</b> class from TIBCO.
<b>Out Total</b>	The total number of outbound messages for the queue.
<b>Pend Msgs</b>	The number of currently pending messages for the queue.

<b>Pend Size</b>	The amount of space, in bytes, used by pending messages for the queue.
<b>activeDurableCount</b>	The current number of active durables.
<b>consumerCount</b>	The number of active and inactive consumers.
<b>durableCount</b>	The number of active and inactive durables.
<b>failSafe</b>	When checked, the message is marked as failsafe delivery.
<b>fcMaxBytes</b>	The maximum number of bytes allocated for use by flow control.
<b>global</b>	When checked, the message is global and is routed to other servers.
<b>inboundByteRate</b>	The amount of inbound messages for the queue, in bytes per second.
<b>inboundTotalBytes</b>	The total amount of inbound messages for the queue, in bytes.
<b>maxBytes</b>	The maximum amount of bytes allocated for use by the queue.
<b>maxMsgs</b>	The maximum number of messages allocated for use by the queue.
<b>outboundByteRate</b>	The amount of outbound messages for the queue, in bytes per second.
<b>outboundTotalBytes</b>	The total amount of outbound messages for the queue, in bytes.
<b>overflowPolicy</b>	Indicates whether an overflow policy is set for the queue: <b>0</b> = No policy is set. <b>1</b> = A policy is set.
<b>secure</b>	When checked, the queue is designated as secure and enforces permission policies.
<b>static</b>	When checked, the queue has a static destination.
<b>subscriberCount</b>	The number of subscribers that receive queue message.
<b>description</b>	Descriptive text to help the administrator identify this resource.
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > <b>(Project Name)</b> > <b>Solution Package Configuration</b> > <b>TIBCO Enterprise Message Service</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.

<b>time_stamp</b>	The date and time this row of data was last updated.
<b>DeltainboundTotalMessages</b>	The change in total inbound messages since the last update.
<b>DeltainboundTotalBytes</b>	The change in total inbound message bytes since the last update.
<b>DeltaoutboundTotalMessages</b>	The change in total outbound messages since the last update.
<b>DeltaoutboundTotalBytes</b>	The change in total outbound message bytes since the last update.
<b>prefetch</b>	Lists the maximum number of messages consumers can fetch.
<b>expiryOverride</b>	If set to a non-zero value for a destination and the server delivers a message to the destination, the server replaces the producer's expiration value with this value.
<b>store</b>	Provides the store for this destination where persistent messages are stored.
<b>deliveredMessageCount</b>	Indicates the total number of messages that have been delivered and acknowledged.
<b>URLQueue</b>	The IP address and port for the queue.
<b>exclusive</b>	When checked, the server sends all messages on this queue to one consumer.
<b>maxRedelivery</b>	The maximum number of attempts for attempting redelivery of a message.
<b>receiverCount</b>	The number of receivers that receive queue message.

## All Queues Summary

Track performance and utilization metrics and trends for all queues on a single server.



**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** , **Table** open commonly accessed displays.
- 6,047** The number of items currently in the display.
- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

**Note:** Clicking **Heatmap** in the Title Bar takes you to the "All Queues Heatmap" display. Clicking **Table** in the Title Bar takes you to the "All Queues Table" display.

**Fields and Data**

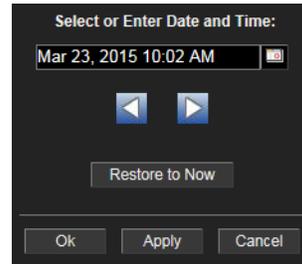
This display includes:

- EMS Server** The EMS Server selected from this drop-down menu populates all associated queue data in this display.
- Name** The name of the server selected in the **EMS Server** drop down list.
- Totals For Server** Shows metrics for all queues on the selected server.

	<p><b>Messages In</b> <b>Msgs/sec</b> -- The total number of inbound messages for all queues on the server, per second.</p> <p><b>Total</b> -- The total number of inbound messages for all queues on the server since the server was started.</p> <p><b>Bytes/sec</b> -- The amount of inbound messages, in bytes per second, for all queues on the server.</p> <p><b>Total</b> -- The amount of inbound messages, in kilobytes, for all queues on the server since the server was started.</p>
	<p><b>Messages Out</b> <b>Msgs/sec</b> -- The total number of outbound messages for all queues on the server, per second.</p> <p><b>Total</b> -- The total number of outbound messages for all queues on the server since the server was started.</p> <p><b>Bytes/sec</b> -- The amount of outbound messages, in bytes per second, for all queues on the server.</p> <p><b>Total</b> -- The amount of outbound messages for all queues on the server, in kilobytes, since the server was started.</p>
	<p><b>Pending Messages</b> <b>Current</b> -- The total number of messages currently waiting to be processed.</p> <p><b>Bytes</b> -- The amount of messages, in bytes, currently waiting to be processed.</p>
<b>Trend Graphs</b>	<p>Shows metrics for all queues on the selected server.</p> <p><b>Pending Msgs</b> -- Traces the number of messages currently waiting to be processed.</p> <p><b>In Msgs / sec</b> -- Traces the number of inbound messages for all queues, per second.</p> <p><b>Out Msgs / sec</b> -- Traces the number of outbound messages for all queues, per second.</p>
	<p><b>Log Scale</b> This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.</p>

**Base at Zero** When this option is checked, zero is set as the Y axis minimum for all graph traces.

**Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



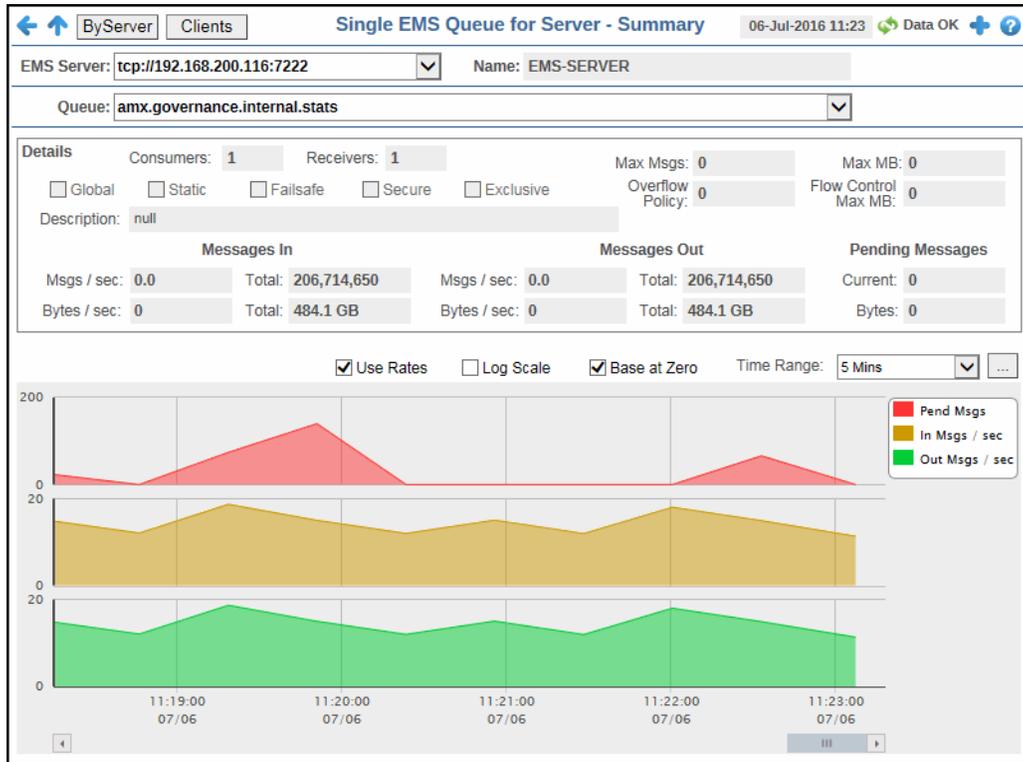
By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Single Queue Summary

Track performance and utilization metrics for a single queue on a single server.



**Title Bar (possible features are):**

- Open the previous and upper display.
- Open the online help page for this display.
- Menu** open commonly accessed displays.
- 6,047** The number of items currently in the display.
- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

**Note:** Clicking **Clients** in the Title Bar takes you to the "Single EMS Queue-Clients" display. Clicking **By Server** in the Title Bar takes you to the "Single Queue By Server".

### Fields and Data

This display includes:

**EMS Server** The EMS Server selected from this drop-down menu populates the **Queues** drop-down menu with the queues belonging to this EMS Server.

**Name** The name of the EMS Server selected from the **EMS Server** drop-down menu.

- Queue** Select a queue from the drop-down menu. The selection made here populates this display.
- Browse** Click to browse the contents of the selected queue in a separate window. The queue browser table displays up to 100,000 rows of messages.

Message ID	Size	Priority	Destination
Msg 1748	908	785	610
Msg 1749	824	852	315
Msg 1753	344	324	315
Msg 1754	334	301	440
Msg 1755	873	918	315
Msg 1756	405	167	558
Msg 1757	791	491	464
Msg 1759	523	499	40
Msg 1758	461	364	640
Msg 1757	684	204	432
Msg 1767	353	461	872
Msg 1758	668	275	611
Msg 1760	592	229	410

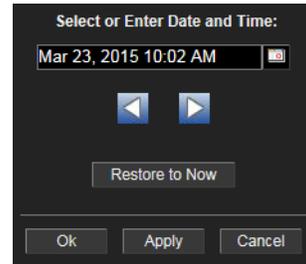
By default, this button is disabled due to the fact that use of this option could significantly impact performance. To enable it, add the following substitution to the properties file with which you execute the Display Server and/or Viewer:

**sl.rtvview.sub=\$emsDestBrowseButtonVisFlag:1**

- Details** Shows metrics for the queue selected from the **Queue** drop-down menu.
- Consumers** The number of consumers currently interacting with the queue.
- Receivers** The number of consumers currently receiving messages from the queue.
- Max Msgs** The maximum number of messages allocated for the queue.
- Max MB** The maximum amount of memory, in megabytes, allocated for use by the queue.
- Global** When checked, the message is global and is routed to other servers.
- Static** When checked, the queue has a static destination.
- Failsafe** When checked, the message is marked as failsafe delivery.
- Secure** When checked, the queue is designated as secure and enforces permission policies.
- Exclusive** When checked, the server sends all messages on this queue to one consumer.
- Overflow Policy** Indicates whether an overflow policy is set for the queue:  
**0** = No policy is set.  
**1** = A policy is set.
- Flow Control Max MB** The maximum amount of memory, in megabytes, allocated for flow control use by the queue.
- Description** Description of the Queue.
- Messages In**
- Msgs/sec** The number of inbound messages, per second, for the selected queue.
- Total** The total number of inbound messages for the selected queue since the server was started.
- Bytes/sec** The size of the inbound messages, in bytes per second, for the selected queue.
- Total** The total size of inbound messages, in bytes, for the selected queue since the server was started.

<b>Messages Out</b>	<b>Msgs/sec</b>	The number of outbound messages, per second, for the selected queue.
	<b>Total</b>	The total number of outbound messages for the selected queue since the server was started.
	<b>Bytes/sec</b>	The size of outbound messages, in bytes per second, for the selected queue.
	<b>Total</b>	The total size of outbound messages, in bytes, for the selected queue since the server was started.
<b>Pending Messages</b>	<b>Current</b>	The total number of messages for the selected queue currently waiting to be processed.
	<b>Bytes</b>	The size, in bytes, of messages for the selected queue currently waiting to be processed.
<b>Trend Graphs</b>	Shows metrics for the selected queue on the specified server.	
	<b>Pending Msgs</b>	-- Traces the number of messages currently waiting to be processed.
	<b>In Msgs / sec</b>	-- Traces the number of inbound messages, per second. This trend graph only displays when <b>Use Rates</b> is selected.
	<b>Out Msgs / sec</b>	-- Traces the number of outbound messages, per second. This trend graph only displays when <b>Use Rates</b> is selected.
	<b>Delta In Msgs</b>	-- Traces the change in total inbound messages since the last update. This trend graph only displays when <b>Use Rates</b> is not selected.
	<b>Delta Out Msgs</b>	-- Traces the change in total inbound messages since the last update. This trend graph only displays when <b>Use Rates</b> is selected.
	<b>Use Rates</b>	When this check box is selected, the inbound and outbound message rates ( <b>In Msgs/sec</b> and <b>Out Msgs/sec</b> ) display in the trend graph. When this check box is not selected, the delta inbound and outbound messages ( <b>Delta In Msgs</b> and <b>Delta Out Msgs</b> ) display in the trend graph.
<b>Log Scale</b>	This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.	

- Base at Zero** When this option is checked, zero is set as the Y axis minimum for all graph traces.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows  to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Single EMS Queue-Clients

View data for all consumers and producers associated with the selected queue.

**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** , **Table** open commonly accessed displays.
- 6,047** The number of items currently in the display.
- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

**Note:** Clicking **By Server** in the Title Bar takes you to the "Single Queue By Server". Clicking **Summary** in the Title Bar takes you to the "Single Queue Summary" display.

### Fields and Data

This display includes:

- EMS Server** The EMS Server selected from this drop-down menu populates the Queue drop-down menu with the Queues belonging to this EMS Server.
- Name** The name of the EMS Server selected from the EMS Server drop-down menu.

<b>Queue</b>	Select a Queue from the drop-down menu to view details for the selected Queue.
<b>Show Active Only</b>	Select this check box to view only the active producers and consumers for the selected EMS Queue.
<b>Producers</b>	Shows data for all producers for the selected queue.
<b>ID</b>	A unique string identifier assigned to each producer.
<b>clientID</b>	A unique string identifier assigned to each client.
<b>Msgs / sec</b>	The number of messages, per second, that are emitted by the producer.
<b>Msgs Total</b>	The total number of messages emitted by the producer since the server was started.
<b>Bytes / sec</b>	The size of messages, in bytes per second, that are emitted by the producer.
<b>Total Bytes</b>	The total size of messages, in bytes, emitted by the producer since the server was started.
<b>userName</b>	The user name.
<b>host</b>	The name of the host.
<b>sessionID</b>	A unique string identifier assigned to each session.
<b>connection ID</b>	A unique string identifier assigned to each connection.
<b>createTime</b>	The amount of time, in milliseconds, since the producer was created.
<b>time_stamp</b>	The date and time this row of data was last updated.
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > <b>(Project Name) &gt; Solution Package Configuration &gt; TIBCO Enterprise Message Service &gt; DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Consumers</b>	Shows data for all consumers associated with the selected queue.
<b>ID</b>	A unique string identifier assigned to each consumer.
<b>clientID</b>	A unique string identifier assigned to each client.
<b>Msgs / sec</b>	The number of messages, per second, that are processed by the consumer.
<b>Msgs Total</b>	The total number of messages that have been processed by the consumer.
<b>Bytes / sec</b>	The size of messages, in bytes per second, that are processed by the consumer.
<b>Total Bytes</b>	The total size of messages, in bytes, processed by the consumer since the server was started.
<b>userName</b>	The user name.
<b>host</b>	The name of the host machine.

<b>Msgs Sentt</b>	<p>The number of messages sent to the consumer that were not yet acknowledged by the consumer's session.</p> <p>The <b>sl.rtvew.jmsadm.queryCIDetails</b> property must be set to <b>true</b> in your <b>sample.properties</b> file to see this column.</p>
<b>Size Msg Sent</b>	<p>The combined size of messages sent to the consumer that were not yet acknowledged by the consumer's session.</p> <p><b>The sl.rtvew.jmsadm.queryCIDetails</b> property must be set to <b>true</b> in your <b>sample.properties</b> file to see this column.</p>
<b>Ack Msgs</b>	<p>The total number of messages that have been sent to the consumer and have been acknowledged by the consumer's session.</p> <p>The <b>sl.rtvew.jmsadm.queryCIDetails</b> property must be set to <b>true</b> in your <b>sample.properties</b> file to see this column.</p>
<b>Sent Msgs</b>	<p>The total number of messages sent to the consumer since the consumer was created.</p> <p>The <b>sl.rtvew.jmsadm.queryCIDetails</b> property must be set to <b>true</b> in your <b>sample.properties</b> file to see this column.</p>
<b>Elap. Since Last Ack</b>	<p>The amount of time (in milliseconds) that has elapsed since the last time a message sent to the consumer was acknowledged by the consumer's session.</p> <p>The <b>sl.rtvew.jmsadm.queryCIDetails</b> property must be set to <b>true</b> in your <b>sample.properties</b> file to see this column.</p>
<b>Elap. Since Last Sent</b>	<p>The amount of time (in milliseconds) that has elapsed since the last time the server sent a message to the consumer.</p> <p>The <b>sl.rtvew.jmsadm.queryCIDetails</b> property must be set to <b>true</b> in your <b>sample.properties</b> file to see this column.</p>
<b>destination Prefetch</b>	<p>The actual destination prefetch value used by the server at runtime.</p> <p>The <b>sl.rtvew.jmsadm.queryCIDetails</b> property must be set to <b>true</b> in your <b>sample.properties</b> file to see this column.</p>
<b>prefetch Delivered Count</b>	<p>The number of prefetch messages delivered to the consumer by the server. For consumers receiving messages on any destination with positive prefetch value, this value is never more than the prefetch value of the destination. This value cannot be used to identify the status of the consumer, but it can be used in conjunction with other consumer information values to identify consumers who stopped receiving messages due to application-specific problems.</p> <p>The <b>sl.rtvew.jmsadm.queryCIDetails</b> property must be set to <b>true</b> in your <b>sample.properties</b> file to see this column.</p>
<b>durable Name</b>	<p>The name of the durable.</p>
<b>routeName</b>	<p>The queue owner server name if the consumer's destination is a routed queue.</p> <p>The <b>sl.rtvew.jmsadm.queryCIDetails</b> property must be set to <b>true</b> in your <b>sample.properties</b> file to see this column.</p>
<b>isActive</b>	<p>When checked, the consumer is active and can receive messages from the server.</p> <p>The <b>sl.rtvew.jmsadm.queryCIDetails</b> property must be set to <b>true</b> in your <b>sample.properties</b> file to see this column.</p>
<b>isSystem</b>	<p>This check box is checked if the consumer was automatically created by the system.</p> <p>The <b>sl.rtvew.jmsadm.queryCIDetails</b> property must be set to <b>true</b> in your <b>sample.properties</b> file to see this column.</p>

<b>sessionAck Mode</b>	Lists the consumer's session acknowledge mode as a constant defined in <b>TibjmsAdmin</b> . The <b>sl.rtvew.jmsadm.queryCIDetails</b> property must be set to <b>true</b> in your <b>sample.properties</b> file to see this column.
<b>session ID</b>	A unique string identifier assigned to each session.
<b>connection ID</b>	A unique string identifier assigned to each connection.
<b>createTime</b>	The amount of time, in milliseconds, since the consumer was created.
<b>time_stamp</b>	The date and time this row of data was last updated.
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > <b>(Project Name) &gt; Solution Package Configuration &gt; TIBCO Enterprise Message Service &gt; DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.

## Single Queue By Server

Track performance and utilization metrics of a single queue across all servers. Compare queue activity among servers.

The screenshot shows the 'EMS Queue - Detail by Server' window. The queue is 'amx.governance.stats'. The table below displays the following data:

URL	Consumers	exclusive	failsafe	fcMaxBytes	global	In Byte Rate	In Msg R
tcp://192.168.200.115:7222	0	<input type="checkbox"/>	<input type="checkbox"/>	0	<input type="checkbox"/>	0	
tcp://192.168.200.116:7222	1	<input type="checkbox"/>	<input type="checkbox"/>	0	<input type="checkbox"/>	17,609	
tcp://192.168.200.121:7222	1	<input type="checkbox"/>	<input type="checkbox"/>	0	<input type="checkbox"/>	0	

**Title Bar** (possible features are):

-   Open the previous and upper display.
-  Open an instance of this display in a new window.
-  Open the online help page for this display.
- ,  open commonly accessed displays.
- The number of items currently in the display.

 **Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

 Open the **Alert Views - RTView Alerts Table** display.

---

**Note:** Clicking **Summary** in the Title Bar takes you to the ["Single Queue Summary"](#). Clicking **Clients** in the Title Bar takes you to the ["Single EMS Queue-Clients"](#) display.

---

### Fields and Data

This display includes:

<b>Queue</b>	The Queue selected from this drop-down menu populates this display.
<b>Table</b>	Shows details about the selected Queue for each server that has the queue defined. Select a server to view details in the <a href="#">"Single Queue Summary"</a> display.
<b>URL</b>	The URL of the server.
<b>Consumers</b>	The number of active and inactive consumers.
<b>exclusive</b>	When checked, the server sends all messages on this queue to one consumer.
<b>failSafe</b>	When checked, the message is marked as failsafe delivery.
<b>fcMaxBytes</b>	The maximum number of bytes allocated for use by flow control.
<b>global</b>	When checked, the message is global and is routed to other servers.
<b>In Byte Rate</b>	The amount of inbound messages for the queue, in bytes per second.
<b>In Msg Rate</b>	The amount of inbound messages for the queue, in number of messages per second.
<b>In Total Bytes</b>	The total number of inbound bytes for the queue.
<b>In Total Msgs</b>	The total number of inbound messages for the queue.
<b>maxBytes</b>	The maximum amount of bytes allocated for use by the queue.
<b>maxMsgs</b>	The maximum number of messages allocated for use by the queue.
<b>maxRedelivery</b>	The maximum number of attempts for attempting redelivery of a message.

<b>Out Byte Rate</b>	The amount of outbound messages (in bytes) per second.
<b>Out Msg Rate</b>	The number of outbound messages per second.
<b>Out Total Bytes</b>	The total amount of outbound messages, in bytes, since the server was started.
<b>Out Total Msgs</b>	The total number of outbound messages since the server was started.
<b>overflowPolicy</b>	Indicates whether an overflow policy is set for the queue: <b>0</b> = No policy is set. <b>1</b> = A policy is set.
<b>Pending Msgs</b>	The number of currently pending messages.
<b>Pending Msgs Size</b>	The amount of space, in bytes, pending messages use for the queue.
<b>Receivers</b>	The number of receivers of queue messages.
<b>secure</b>	When checked, the topic is designated as secure and enforces permission policies.
<b>static</b>	When checked, the topic has a static destination.
<b>time_stamp</b>	The date and time this row of data was last updated.
<b>description</b>	Descriptive text to help the administrator identify this resource.

## EMS Clients

These displays present performance metrics for all server connections, including users, routes between servers, producers, consumers and durables connected to a specific EMS server.

- **"Connections"**: Shows connection information on a single server.
- **"Bridges, Users, Ports"**: Shows utilization metrics for bridges, users and ports on a single server.
- **"Routes"**: Shows bridges for server routes on a single server.
- **"Producers"**: Shows utilization metrics for producers on a single server.
- **"Producer Summary"**: Shows utilization metrics for producers on a single server.
- **"Consumers"**: Shows utilization metrics for consumers on a single server.
- **"Consumer Summary"**: Shows utilization metrics for consumers on a single server.
- **"Durables"**: Shows utilization metrics for durables on a single server.

## Connections

View connections on a single server.

Conn ID	Client ID	Conn URL	User	host	type	consumerCount	producerCount
1181372		[admin@SLHOST16]	admin	SLHOST16	CONN...	3	
1181373	EMSGMS.Unbound...	[admin@SLHOST16]	admin	SLHOST16	CONN...	0	
1181375		[admin@SLHOST16]	admin	SLHOST16	CONN...	3	
1181376	EMSGMS.Unbound...	[admin@SLHOST16]	admin	SLHOST16	CONN...	0	
1181378		[admin@SLHOST16]	admin	SLHOST16	CONN...	3	
1181379	EMSGMS.amxadmi...	[admin@SLHOST16]	admin	SLHOST16	CONN...	0	
1181388		[admin@SLHOST16]	admin	SLHOST16	CONN...	0	
1181396		[admin@SLHOST16]	admin	SLHOST16	CONN...	0	
1181398		[admin@SLHOST16]	admin	SLHOST16	CONN...	3	
1181399	EMSGMS.amxadmi...	[admin@SLHOST16]	admin	SLHOST16	CONN...	0	
1181402		[admin@SLHOST16]	admin	SLHOST16	CONN...	0	
1181403		[admin@SLHOST16]	admin	SLHOST16	CONN...	3	
1181405	EMSGMS.amxadmi...	[admin@SLHOST16]	admin	SLHOST16	CONN...	0	
1181407		[admin@SLHOST16]	admin	SLHOST16	QUEUE	0	
1181410		[admin@SLHOST16]	admin	SLHOST16	CONN...	0	
1181415		[admin@SLHOST16]	admin	SLHOST16	QUEUE	0	
1181424	AMX_MGMT.DevEn...	[admin@SLHOST16]	admin	SLHOST16	CONN...	4	
1181427	AMX_SV:8bf3d299-...	[admin@SLHOST16]	admin	SLHOST16	CONN...	37	
1181590		[admin@SLHOST16]	admin	SLHOST16	CONN...	3	
1181591	EMSGMS._amxadm...	[admin@SLHOST16]	admin	SLHOST16	CONN...	0	
1181593		[admin@SLHOST16]	admin	SLHOST16	CONN...	3	
1181594	EMSGMS.Unbound...	[admin@SLHOST16]	admin	SLHOST16	CONN...	0	
1181596		[admin@SLHOST16]	admin	SLHOST16	CONN...	3	
1181597	EMSGMS.Unbound...	[admin@SLHOST16]	admin	SLHOST16	CONN...	0	
1181599		[admin@SLHOST16]	admin	SLHOST16	CONN...	3	
1181600	EMSGMS.amxadmi...	[admin@SLHOST16]	admin	SLHOST16	CONN...	0	
1181603		[admin@SLHOST16]	admin	SLHOST16	CONN...	7	
1181609	AMX_MGMT.Syste...	[admin@SLHOST16]	admin	SLHOST16	CONN...	5	

### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

### Fields and Data

This display includes:

- EMS Server** The EMS Server selected from this drop-down menu populates all associated Connections data in this display.
- Name** The name of the EMS Server selected from the EMS Server drop-down menu.
- Show Active Only** Select this check box to display only active connections.
- Client ID Filter** Filter field that allows you to filter the list of connections by client ID.

**Filtered Connection Count**

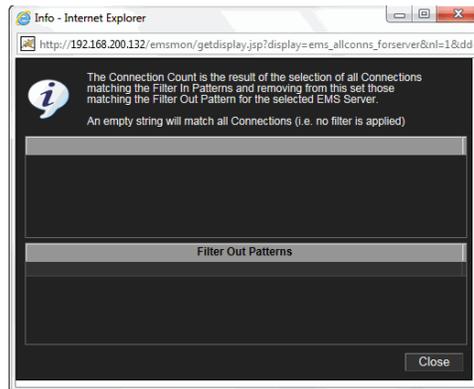
This field is broken into two different values. The first value is the total number of currently active connections on the selected server, which is filtered by the **Filter Pattern** field and by the default value specified in the **\$emsConnectionFilterOutPattern** property in the **emsmon/conf/rtvapl.properties** file. The second value is the total number of connections on the selected server. In other words, the filtered number of connections/the total number of connections on the server.

The default value for the **\$emsConnectionFilterOutPattern** property is:

```
collector.sl.rtvapl.sub=$emsConnectionFilterOutPattern: '^(?!^\\[admin\\@)'
```

You can modify the filter value by editing the **\$emsConnectionFilterOutPattern** property in the **sample.properties File**, which will override the default value.

Clicking the associated Help button  displays the **Info** dialog, which displays the defined filter in and filter out properties used by the **Filtered Connection Count**.



**User Filter** Filter field that allows you to filter the list of connections by user name.

**Connections** This table describes the current connections on the selected server.

<b>Conn ID</b>	The unique numeric ID assigned to this connection that can be used for deletion.
<b>Client ID</b>	The unique string identifier assigned to the client.
<b>Conn URL</b>	The connection URL.
<b>User</b>	The user name.
<b>host</b>	The name of the host to which the server is connected.
<b>type</b>	The type of connection: Queue, Topic or System.
<b>consumerCount</b>	The total number of consumers currently connected.
<b>producerCount</b>	The total number of producers currently connected.
<b>sessionCount</b>	The total number of sessions currently connected.
<b>startTime</b>	The date and time the server was started
<b>upTime</b>	The amount of time, in milliseconds, since the server was started.

**Expired**

When checked, performance data has not been received within the time specified (in seconds) in the **Expire Time** field in the **Duration** region in the RTView Configuration Application > **(Project Name)** > **Solution Package Configuration** > **TIBCO Enterprise Message Service** > **DATA STORAGE** tab. The **Delete Time** field (also in the **Duration** region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.

**time\_stamp**

The date and time this row of data was last updated.

**Bridges, Users, Ports**

View bridges configured on an EMS Server, as well as their associated users and ports. You can right-click in the **Bridges** table and select **Go To Source** to view bridged source information in the "Single Queue Summary" if the source is a queue, or "Single Topic Summary" if the source is a topic. You can right-click in the **Bridges** table and select **Go To Target** to view bridged target information in the "Single Queue Summary" if the target is a queue, or "Single Topic Summary" if the target is a topic.

---

**Note:** The **Go To Source** option will not be enabled if the source side of the bridge is wildcarded.

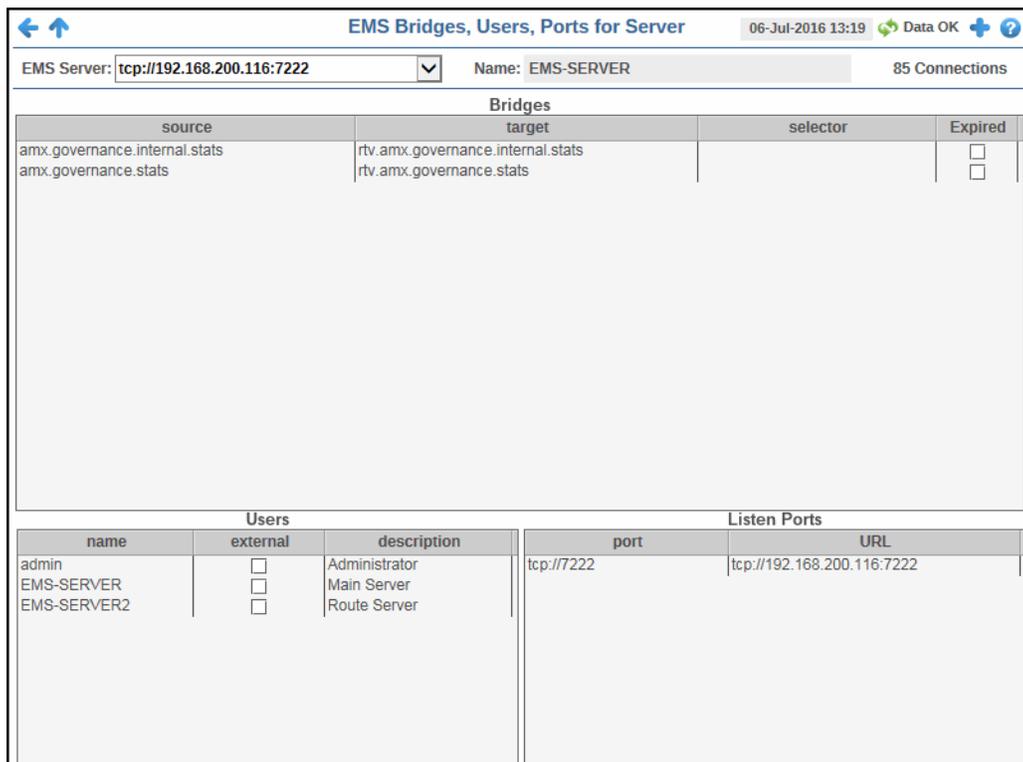
---



---

**Note:** The functionality of the **Drop Down** option in the drop down list that displays when you right-click in the **Bridges** table is replaced by the **Go To Source** and **Go To Target** options, and no additional functionality exists for the **Drop Down** option.

---



**Title Bar** (possible features are):

-   Open the previous and upper display.
-  Open an instance of this display in a new window.
-  Open the online help page for this display.
- ,  open commonly accessed displays.
- The number of items currently in the display.

 **Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

 Open the **Alert Views - RTView Alerts Table** display.

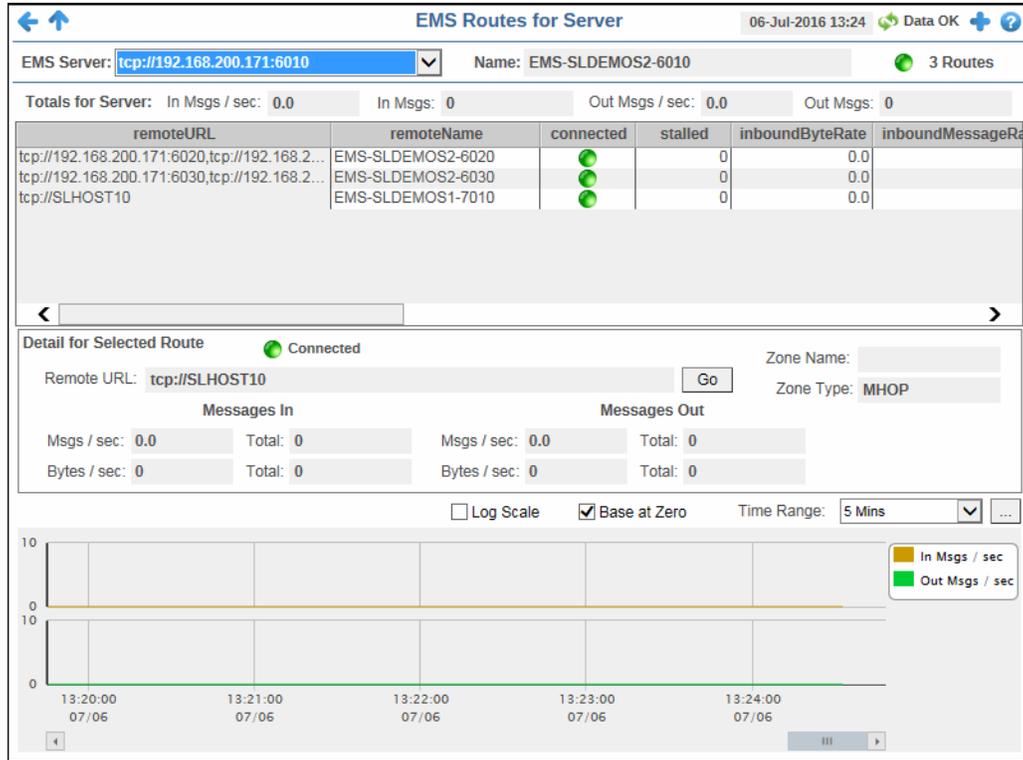
## Fields and Data

This display includes:

<b>EMS Server</b>	The EMS Server selected from this drop-down menu populates all associated Bridges, Users, and Ports data in this display.								
<b>Name</b>	The name of the EMS Server selected from the <b>EMS Server</b> drop-down menu.								
<b>Bridges</b>	This table describes the bridges for the selected server. <table border="0" style="margin-left: 20px;"> <tr> <td style="vertical-align: top;"><b>source</b></td> <td>The topic or queue which is the source of the bridge.</td> </tr> <tr> <td style="vertical-align: top;"><b>target</b></td> <td>The topic or queue which is the target of the bridge.</td> </tr> <tr> <td style="vertical-align: top;"><b>selector</b></td> <td>The message selector string or blank if none has been set.</td> </tr> <tr> <td style="vertical-align: top;"><b>Expired</b></td> <td>When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application &gt; <b>(Project Name)</b> &gt; <b>Solution Package Configuration</b> &gt; <b>TIBCO Enterprise Message Service</b> &gt; <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.</td> </tr> </table>	<b>source</b>	The topic or queue which is the source of the bridge.	<b>target</b>	The topic or queue which is the target of the bridge.	<b>selector</b>	The message selector string or blank if none has been set.	<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > <b>(Project Name)</b> > <b>Solution Package Configuration</b> > <b>TIBCO Enterprise Message Service</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>source</b>	The topic or queue which is the source of the bridge.								
<b>target</b>	The topic or queue which is the target of the bridge.								
<b>selector</b>	The message selector string or blank if none has been set.								
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > <b>(Project Name)</b> > <b>Solution Package Configuration</b> > <b>TIBCO Enterprise Message Service</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.								
<b>Users</b>	This table describes the users on the selected server. <table border="0" style="margin-left: 20px;"> <tr> <td style="vertical-align: top;"><b>name</b></td> <td>The name of the connected user.</td> </tr> <tr> <td style="vertical-align: top;"><b>external</b></td> <td>When checked, the user is defined in an external system.</td> </tr> <tr> <td style="vertical-align: top;"><b>description</b></td> <td>Textual description of the user.</td> </tr> </table>	<b>name</b>	The name of the connected user.	<b>external</b>	When checked, the user is defined in an external system.	<b>description</b>	Textual description of the user.		
<b>name</b>	The name of the connected user.								
<b>external</b>	When checked, the user is defined in an external system.								
<b>description</b>	Textual description of the user.								
<b>Listen Ports</b>	This table describes the connections the selected server is to listen for. <table border="0" style="margin-left: 20px;"> <tr> <td style="vertical-align: top;"><b>port</b></td> <td>The IP address and port number on which the server is to listen for connections.</td> </tr> <tr> <td style="vertical-align: top;"><b>URL</b></td> <td>The URL on which the server is to listen for connections.</td> </tr> </table>	<b>port</b>	The IP address and port number on which the server is to listen for connections.	<b>URL</b>	The URL on which the server is to listen for connections.				
<b>port</b>	The IP address and port number on which the server is to listen for connections.								
<b>URL</b>	The URL on which the server is to listen for connections.								

## Routes

Track utilization metrics for server routes on a single server. Inbound metrics, such as **inboundByteRate**, indicate an in route to the server. Outbound metrics, such as **outboundByteRate**, indicate an out route to the server.



**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

### Fields and Data

This display includes:

- EMS Server** The EMS Server selected from this drop-down menu populates all associated Routes data in this display.
- Name** The name of the EMS server selected from the **EMS Server** drop-down menu.
- Routes** The number of server routes and the connection state.

-  -- One or more routes for this server are disconnected.
-  -- All routes for this server are connected.
-  -- There are no routes for this server.

**Totals For Server**

Shows metrics for all server routes on the selected server.

<b>In Msgs / sec</b>	The number of inbound messages, per second.
<b>In Msgs</b>	The total number of inbound messages.
<b>Out Msgs / sec</b>	The number of outbound messages, per second.
<b>Out Msgs</b>	The total number of outbound messages.

**Table**

This table shows metrics for each server route on the selected server. Select a route to view details.

<b>remoteURL</b>	The URL of the remote server.
<b>remoteName</b>	The name of the remote server.
<b>connected</b>	When checked, the server route is connected.
<b>stalled</b>	Indicates whether the IO flow stalled on the route. A value of <b>0</b> (zero) = not stalled. A value of <b>1</b> = stalled.
<b>inboundByteRate</b>	The rate of inbound data in bytes, per second.
<b>inboundMessageRate</b>	The rate of inbound messages in number of messages per second.
<b>inboundTotalBytes</b>	The total number of inbound bytes.
<b>inboundTotalMessages</b>	The total number of inbound messages.
<b>outboundByteRate</b>	The rate of outbound data in bytes per second.
<b>outboundMessageRate</b>	The rate of outbound messages in number of messages per second.
<b>outboundTotalBytes</b>	The total number of outbound bytes.
<b>outboundTotalMessages</b>	The total number of outbound messages.
<b>zoneName</b>	The name of the zone for the route.
<b>zoneType</b>	Indicates a multi-hop or one-hop zone.
<b>active</b>	Indicates whether the server route is currently transferring data: <b>1</b> = true <b>0</b> = false

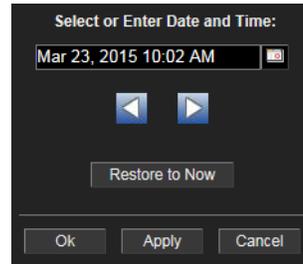
	<b>inactive</b>	Indicates whether the server route is currently transferring data: <b>1</b> = true <b>0</b> = false	
	<b>suspended</b>	Indicates whether outbound messages to the route have been suspended: <b>1</b> = true <b>0</b> = false	
	<b>remoteURLName</b>	The IP address and name for the remote connection.	
<b>Detail for Selected Route</b>	Shows metrics for the server route selected from the table.		
	<b>Connected</b>	The server route connection state.  -- The server route is disconnected  -- The server route is connected.	
	<b>Zone Name</b>	The name of the zone for the route.	
	<b>Remote URL</b>	The IP address and port number for the server route connection. Click the  button to open the selected server in the EMS Single Server Summary display.	
	<b>Zone Type</b>	Indicates a multi-hop or one-hop zone.	
	<b>Messages In</b>	<b>Msgs/sec</b> -- The number of inbound messages, per second. <b>Total</b> -- The total number of inbound messages since the connection was established. <b>Bytes/sec</b> -- The amount of inbound messages, in bytes per second, for this server route. <b>Total</b> -- The amount of inbound messages, in kilobytes, for this server route since the connection was established.	
	<b>Messages Out</b>	<b>Msgs/sec</b> -- The number of outbound messages, per second. <b>Total</b> -- The total number of outbound messages since the connection was established. <b>Bytes/sec</b> -- The amount of outbound messages, in bytes per second. <b>Total</b> -- The amount of outbound messages, in kilobytes, since the connection was established.	
	<b>Trend Graphs</b>	Shows message data for the selected route.	
		<b>In Msgs / sec</b>	-- Traces the number of inbound messages, per second.
		<b>Out Msgs / sec</b>	-- Traces the number of outbound messages, per second.
	<b>Log Scale</b>	This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.	

**Base at Zero**

When this option is checked, zero is set as the Y axis minimum for all graph traces.

**Time Range**

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows  to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Producers

Track utilization metrics for producers on a single server.

ID	Client ID	Destination	Msgs / sec	Total Msgs	By
2007		\$sys.admin	0.0	3,302,620	
2049		\$sys.admin	0.0	2,196,187	
2055		\$sys.admin	0.0	276,789	
2121		\$sys.admin	0.0	348,891	
2136		\$sys.admin	4.0	43,490	
2138		\$sys.admin	0.0	22,003	
2139		\$sys.admin	0.0	11,948	
2140		\$sys.admin	0.0	11,229	

### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** **Table** open commonly accessed displays.
- 6,047** The number of items currently in the display.

- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

**Note:** Clicking on a row in the Producers table and then clicking the Dest. button in the Title Bar takes you to the "Single Queue Summary" display for the selected producer.

### Fields and Data

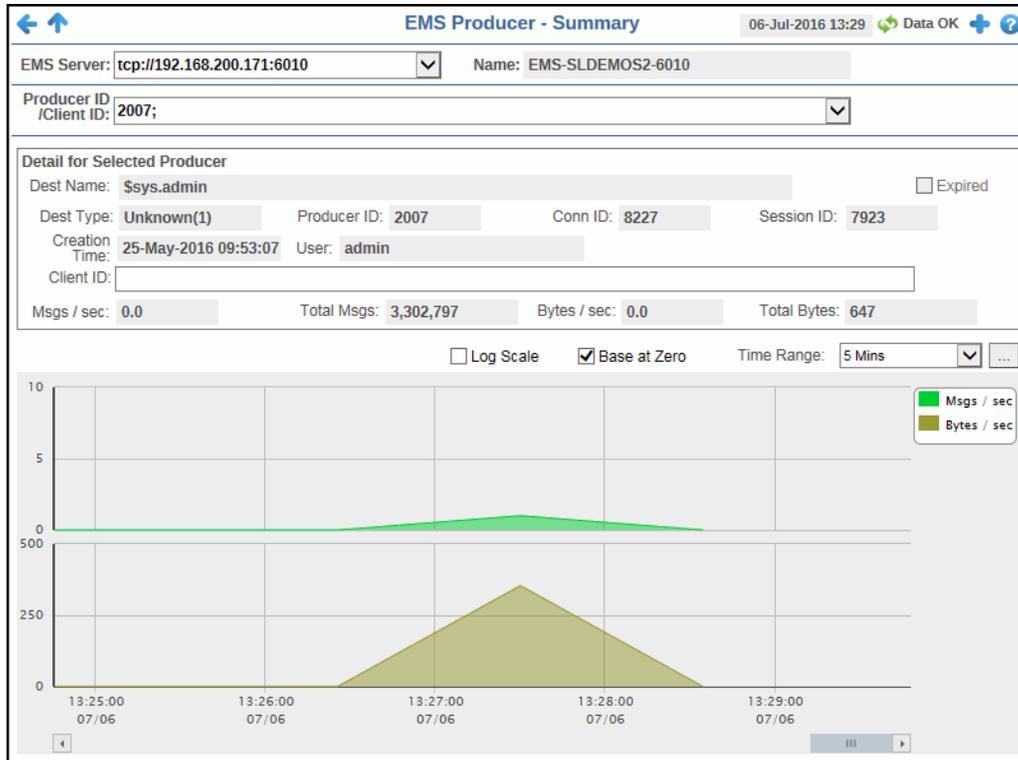
This display includes:

- EMS Server** The EMS Server selected from this drop-down list displays a list of the currently connected Producers.
- Name** The name of the EMS server selected from the **EMS Server** drop-down menu.

<b>Producers</b>	The number of currently connected producers on the server.
<b>Client ID Filter</b>	Filter field that allows you to filter the list of producers by client ID.
<b>DestName Filter</b>	Filter field that allows you to filter the list of producers by destination name.
<b>Show Active Only</b>	Select this check box to display only active producers.
<b>Count</b>	The number of currently connected producers on the server.
<b>Msgs / sec</b>	The number of messages, per second, for the producer.
<b>Total Msgs</b>	The total number of messages for the producer.
<b>Bytes / sec</b>	The amount of messages, in bytes per second, for the producer.
<b>Total Bytes</b>	The total size of messages, in bytes, for the producer.
<b>Table</b>	This table shows metrics for each producer on the selected server. Double-clicking on a row in the Producers table displays details for the producer in the <a href="#">"Producer Summary"</a> drill-down display.
<b>ID</b>	A unique string identifier assigned to each producer.
<b>Client ID</b>	A unique string identifier assigned to each client.
<b>Destination</b>	The name of the destination.
<b>Msgs / sec</b>	The number of messages, per second, for the producer.
<b>Total Msgs</b>	The total number of messages for the producer.
<b>Bytes / sec</b>	The size of messages, in bytes per second, for the producer.
<b>Total Bytes</b>	The total size of messages, in bytes, for the producer.
<b>User</b>	The user name.
<b>Host</b>	The name of the host.
<b>sessionID</b>	A unique string identifier assigned to each session.
<b>ConnID</b>	A unique string identifier assigned to each connection.
<b>createTime</b>	The amount of time, in milliseconds, since the producer was created.
<b>time_stamp</b>	The date and time this row of data was last updated.
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > ( <b>Project Name</b> ) > <b>Solution Package Configuration</b> > <b>TIBCO Enterprise Message Service</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>destinationType</b>	The configured destination type.

## Producer Summary

Displays details for an individual producer. You can access this display by double-clicking on a producer in the "Producers" display.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- 6,047** The number of items currently in the display.

- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

### Fields and Data

This display includes:

- EMS Server** The selected EMS Server populates the Producer ID/ Client ID drop-down menu with associated Producer IDs/Client IDs. This drop down list defaults to the EMS Server that was selected on the previous display.
- Name** The name of the EMS server selected from the **EMS Server** drop-down menu.
- Producer ID/ Client ID** Drop-down menu containing the Producer IDs/Client IDs. This drop down list defaults to the Producer ID/Client ID that was selected on the previous display.

**Detail for Selected Producer**

Shows metrics for the producer selected from the table.

<b>Dest Name</b>	The name of the destination.
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > ( <b>Project Name</b> ) > <b>Solution Package Configuration</b> > <b>TIBCO Enterprise Message Service</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Dest Type</b>	The configured destination type.
<b>Producer ID</b>	A unique string identifier assigned to each producer.
<b>Conn ID</b>	A unique string identifier assigned to each connection.
<b>Session ID</b>	A unique string identifier assigned to each session.
<b>Creation Time</b>	The amount of time, in milliseconds, since the producer was created.
<b>User</b>	The user name.
<b>Client ID</b>	A unique string identifier assigned to each client.
<b>Msgs/sec</b>	The number of messages, per second, for the producer.
<b>Total Msgs</b>	The total number of messages for the producer.
<b>Bytes/sec</b>	The size of messages, in bytes per second, for the producer.
<b>Total Bytes</b>	The total size of messages, in bytes, for the producer.

**Trend Graphs**

Shows message data for the selected producer.

**Msgs / sec** -- Traces the number of messages for the producer, per second.

**Bytes / sec** -- Traces the size of messages for the producer, in bytes.

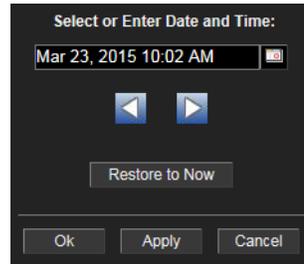
**Log Scale** This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.

**Base at Zero**

When this option is checked, zero is set as the Y axis minimum for all graph traces.

**Time Range**

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Consumers

Track utilization metrics for consumers on a single server.

ID	Client ID	Dest Name	Msgs / sec	Total Msgs	By
7		\$TMP\$.EMS-SLDEMOS2-6020.>	0.0	0	
9		\$TMP\$.EMS-SLDEMOS2-6030.>	0.0	0	
2583		\$TMP\$.EMS-SLDEMOS2-6010.53A25602AAE11EF3...	4.0	3,304,272	
2653		\$TMP\$.EMS-SLDEMOS2-6010.53A25602AAE11F2B...	0.0	2,197,813	
2667		\$TMP\$.EMS-SLDEMOS2-6010.53A25602AAE11F35...	0.0	276,789	
2785		\$TMP\$.EMS-SLDEMOS2-6010.53A25602AAE11F91...	0.0	350,505	
2808		\$TMP\$.EMS-SLDEMOS1-7010.>	0.0	0	
2809		\$TMP\$.EMS-SLDEMOS1-7020.>	0.0	0	
2810		\$TMP\$.EMS-SLDEMOS2-6010.53A25602AAE11FA5.6	4.0	45,130	
2812		\$TMP\$.EMS-SLDEMOS2-6010.53A25602AAE11FA7.4	0.0	23,655	
2813		\$TMP\$.EMS-SLDEMOS2-6010.53A25602AAE11FA8.2	0.0	13,594	
2814		\$TMP\$.EMS-SLDEMOS2-6010.53A25602AAE11FA9...	0.0	12,881	

### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

**Note:** Clicking on a row in the Consumers table and then clicking the **Dest.** button in the Title Bar takes you to the "Single Topic Summary" display for the selected consumer.

### Fields and Data

This display includes:

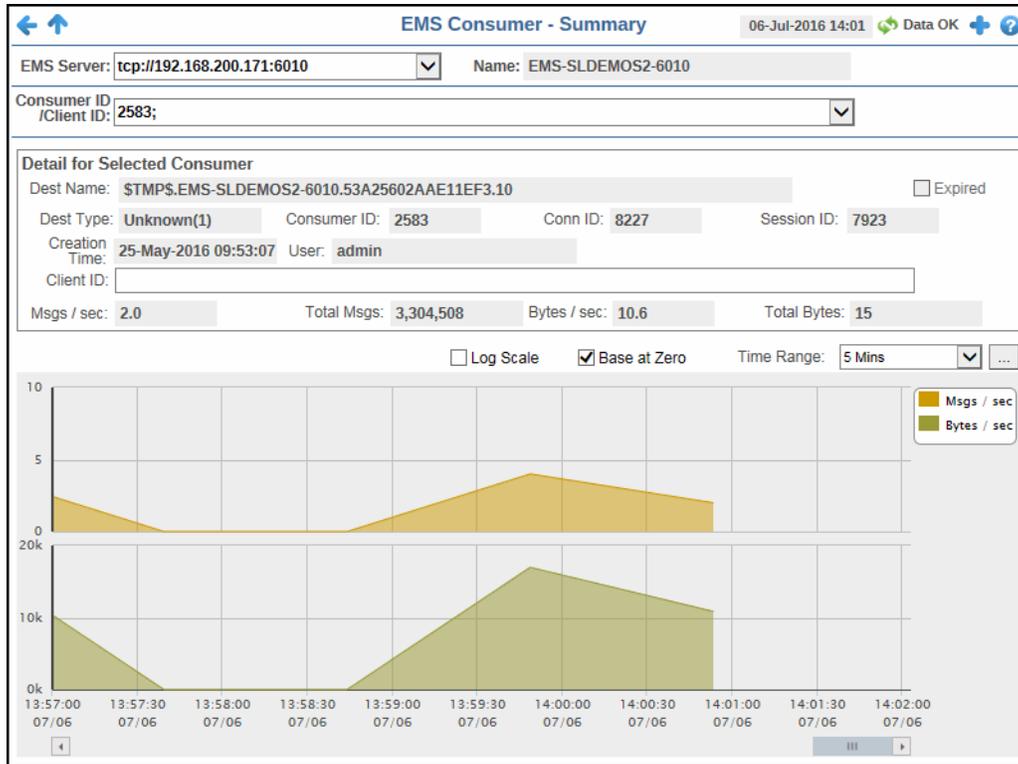
- EMS Server** The EMS Server selected from this drop-down list displays a list of the currently connected Consumers.
- Name** The name of the EMS Server selected from the EMS Server drop-down menu.

<b>Consumers</b>	The number of currently connected consumers on the server.
<b>Client ID Filter</b>	Filter field that allows you to filter the list of consumers by client ID. This filter works in conjunction with the <b>DestName Filter</b> to display the list of consumers.
<b>DestName Filter</b>	Filter field that allows you to filter the list of consumers by destination name. This filter works in conjunction with the <b>Client ID Filter</b> to display the list of consumers.
<b>Show Active Only</b>	Select this check box to display only active consumers.
<b>Count</b>	The number of currently connected consumers on the server.
<b>Msgs / sec</b>	The number of messages, per second, for the consumer.
<b>Bytes / sec</b>	The amount of messages, in bytes per second, for the consumer.
<b>Total Msgs</b>	The total number of messages for the consumer.
<b>Total Bytes</b>	The total size of messages, in bytes, for the consumer.
<b>Table</b>	This table shows metrics for each consumer on the selected server. Double-clicking on a row in the Consumers table displays details for the consumer in the <a href="#">"Consumer Summary"</a> drill-down display.
<b>ID</b>	A unique string identifier assigned to each consumer.
<b>Client ID</b>	A unique string identifier assigned to each client.
<b>Dest Name</b>	The name of the destination.
<b>Msgs / sec</b>	The number of messages, per second, for the consumer.
<b>Total Msgs</b>	The total number of messages for the consumer.
<b>Bytes / sec</b>	The size of messages, in bytes per second, for the consumer.
<b>Total Bytes</b>	The total size of messages, in bytes, for the consumer.
<b>User</b>	The user name.
<b>Host</b>	The name of the host machine.
<b>Session ID</b>	A unique string identifier assigned to each session.
<b>Conn ID</b>	A unique string identifier assigned to each connection.
<b>Curr Msg Sent Count</b>	The number of messages sent to the consumer that were not yet acknowledged by the consumer's session. The <b>sl.rtvview.jmsadm.queryCIDetails</b> property must be set to <b>true</b> in your <b>sample.properties</b> file to see this column.
<b>Curr Msg Sent Size</b>	The combined size of messages sent to the consumer that were not yet acknowledged by the consumer's session. <b>Note:</b> The <b>sl.rtvview.jmsadm.queryCIDetails</b> property must be set to <b>true</b> in your <b>sample.properties</b> file to see this column.
<b>Total Msg Ack Count</b>	The total number of messages that have been sent to the consumer and have been acknowledged by the consumer's session. <b>Note:</b> The <b>sl.rtvview.jmsadm.queryCIDetails</b> property must be set to <b>true</b> in your <b>sample.properties</b> file to see this column.
<b>Total Msg Sent Count</b>	The total number of messages sent to the consumer since the consumer was created. <b>Note:</b> The <b>sl.rtvview.jmsadm.queryCIDetails</b> property must be set to <b>true</b> in your <b>sample.properties</b> file to see this column.

<b>Elapsed Since Last Ack</b>	The amount of time (in milliseconds) that has elapsed since the last time a message sent to the consumer was acknowledged by the consumer's session. <b>Note:</b> The <code>sl.rtvview.jmsadm.queryCIDetails</code> property must be set to <code>true</code> in your <code>sample.properties</code> file to see this column.
<b>Elapsed Since Last Sent</b>	The amount of time (in milliseconds) that has elapsed since the last time the server sent a message to the consumer. <b>Note:</b> The <code>sl.rtvview.jmsadm.queryCIDetails</code> property must be set to <code>true</code> in your <code>sample.properties</code> file to see this column.
<b>Destination Prefetch</b>	The actual destination prefetch value used by the server at runtime. <b>Note:</b> The <code>sl.rtvview.jmsadm.queryCIDetails</code> property must be set to <code>true</code> in your <code>sample.properties</code> file to see this column.
<b>Prefetch Deliv Count</b>	The number of prefetch messages delivered to the consumer by the server. For consumers receiving messages on any destination with positive prefetch value, this value is never more than the prefetch value of the destination. This value cannot be used to identify the status of the consumer, but it can be used in conjunction with other consumer information values to identify consumers who stopped receiving messages due to application-specific problems. <b>Note:</b> The <code>sl.rtvview.jmsadm.queryCIDetails</code> property must be set to <code>true</code> in your <code>sample.properties</code> file to see this column.
<b>Durable Name</b>	The name of the durable.
<b>Route Name</b>	The queue owner server name if the consumer if the consumer's destination is a routed queue. <b>Note:</b> The <code>sl.rtvview.jmsadm.queryCIDetails</code> property must be set to <code>true</code> in your <code>sample.properties</code> file to see this column.
<b>Is Active</b>	When checked, the consumer is active and can receive messages from the server. <b>Note:</b> The <code>sl.rtvview.jmsadm.queryCIDetails</code> property must be set to <code>true</code> in your <code>sample.properties</code> file to see this column.
<b>Is System</b>	This check box is checked if the consumer was automatically created by the system. <b>Note:</b> The <code>sl.rtvview.jmsadm.queryCIDetails</code> property must be set to <code>true</code> in your <code>sample.properties</code> file to see this column.
<b>Session Ack Mode</b>	Lists the consumer's session acknowledge mode as a constant defined in <code>TibjmsAdmin</code> . <b>Note:</b> The <code>sl.rtvview.jmsadm.queryCIDetails</code> property must be set to <code>true</code> in your <code>sample.properties</code> file to see this column.
<b>Create Time</b>	The amount of time, in milliseconds, since the consumer was created.
<b>time_stamp</b>	The date and time this row of data was last updated.
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > ( <b>Project Name</b> ) > <b>Solution Package Configuration</b> > <b>TIBCO Enterprise Message Service</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Dest Type</b>	The configured destination type.

## Consumer Summary

Displays details for an individual consumer. You can access this display by double-clicking on a producer in the “Consumers” display.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- 6,047 The number of items currently in the display.

- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.
- Open the Alert Views - RTView Alerts Table display.

### Fields and Data

This display includes:

- EMS Server** The selected EMS Server populates the Consumer ID/ Client ID drop-down menu with Consumer IDs/Client IDs belonging to this EMS Server. This drop down list defaults to the EMS Server that was selected on the previous display.
- Name** The name of the EMS Server selected from the **EMS Server** drop-down menu.
- Consumer ID/ Client ID** Drop-down menu containing the Consumer IDs/Client IDs. This drop down list defaults to the Consumer ID/Client ID that was selected on the previous display.

**Detail for Selected Consumer**

Shows metrics for the consumer selected from the table.

<b>Dest Name</b>	The name of the destination.
<b>Expired</b>	When checked, performance data has not been received within the time specified (in seconds) in the <b>Expire Time</b> field in the <b>Duration</b> region in the RTView Configuration Application > ( <b>Project Name</b> ) > <b>Solution Package Configuration</b> > <b>TIBCO Enterprise Message Service</b> > <b>DATA STORAGE</b> tab. The <b>Delete Time</b> field (also in the <b>Duration</b> region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
<b>Dest Type</b>	The configured destination type.
<b>Consumer ID</b>	A unique string identifier assigned to each consumer.
<b>Conn ID</b>	A unique string identifier assigned to each connection.
<b>Session ID</b>	A unique string identifier assigned to each session.
<b>Creation Time</b>	The amount of time, in milliseconds, since the consumer was created.
<b>User</b>	The user name.
<b>Client ID</b>	A unique string identifier assigned to each client.
<b>Msgs/sec</b>	The number of messages, per second, for the consumer.
<b>Total Msgs</b>	The total number of messages for the consumer.
<b>Bytes/sec</b>	The size of messages, in bytes per second, for the consumer.
<b>Total Bytes</b>	The total size of messages, in bytes, for the consumer.

**Trend Graphs**

Shows message data for the selected producer.

**Msgs / sec** -- Traces the number of messages for the consumer, per second.

**Bytes / sec** -- Traces the size of messages for the consumer, in bytes.

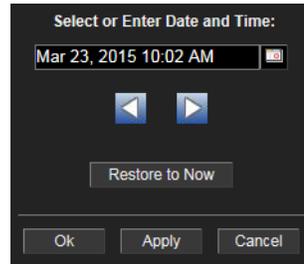
**Log Scale** This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.

**Base at Zero**

When this option is checked, zero is set as the Y axis minimum for all graph traces.

**Time Range**

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Durables

Track utilization metrics for durables on a single server.

### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- 6,047** The number of items currently in the display.

- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

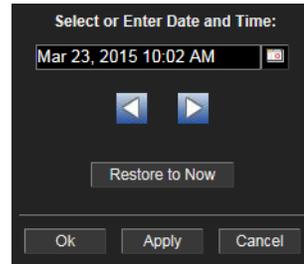
### Fields and Data

This display includes:

- EMS Server** The EMS Server selected from this drop-down menu populates all associated Durables data in this display.
- Name** The name of the EMS Server selected from the **EMS Server** drop-down menu.
- Total Pending Msgs** The total number of pending messages for the durable.
- Durable Count** The number of currently connected durables on the server.
- Table** This table shows metrics for each durable on the selected server.

<b>Name</b>	The name of the durable.
<b>Topic</b>	The name of the topic.
<b>Active</b>	Indicates whether the durable is active.
<b>Client ID</b>	A unique string identifier assigned to each client.
<b>consumerID</b>	A unique string identifier assigned to each consumer.
<b>NoLocalEnabled</b>	Indicates whether the subscriber receives messages from all connections its local connection. <b>Enabled</b> -- The subscriber does not receive messages sent from its local connection. <b>Disabled</b> -- The subscriber receives messages from all connections.
<b>Pending Msgs</b>	The total number of pending messages for the selected durable.
<b>Pending Size</b>	The total amount of pending messages, in bytes, for the selected durable.
<b>Selector</b>	Indicates that the subscriber only receives messages that match this selector.
<b>userName</b>	The name of the user of this durable subscriber.
<b>time_stamp</b>	The date and time this row of data was last updated.
<b>Durable</b>	The name of the durable selected from the table.
<b>Users</b>	The names of the users of this durable subscriber.
<b>Topic</b>	The name of the topic.
<b>Selectors</b>	Indicates that the subscriber only receives messages that match this selector.
<b>Client ID</b>	A unique string identifier assigned to each client.
<b>Consumer ID</b>	A unique string identifier assigned to each consumer.
<b>Pending Msgs</b>	The total number of pending messages for the selected durable.
<b>Pending Msg Size</b>	The total size of pending messages, in bytes, for the selected durable.
<b>Active</b>	Indicates whether the durable is active.
<b>No Local</b>	Indicates whether the subscriber receives messages from all connections its local connection. <b>Enabled</b> The subscriber does not receive messages sent from its local connection. <b>Disabled</b> The subscriber receives messages from all connections.
<b>Trend Graphs</b>	Shows message data for the selected consumer. <b>Pending Msgs</b> -- Traces the number of pending messages for the durable.
<b>Log Scale</b>	This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.

- Base at Zero** When this option is checked, zero is set as the Y axis minimum for all graph traces.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

---

## TIBCO Enterprise Message Service - HTML

The HTML version features an overview display, "**TIBCO EMS Overview - HTML**" (pictured below), and the following Views which can be found under **Components tab > Middleware > TIBCO EMS**:

- "EMS Servers - HTML"
- "EMS Topics - HTML"
- "EMS Queues - HTML"
- "EMS Routes - HTML"
- "EMS Producers - HTML"
- "EMS Consumers - HTML"
- "EMS Durables - HTML"

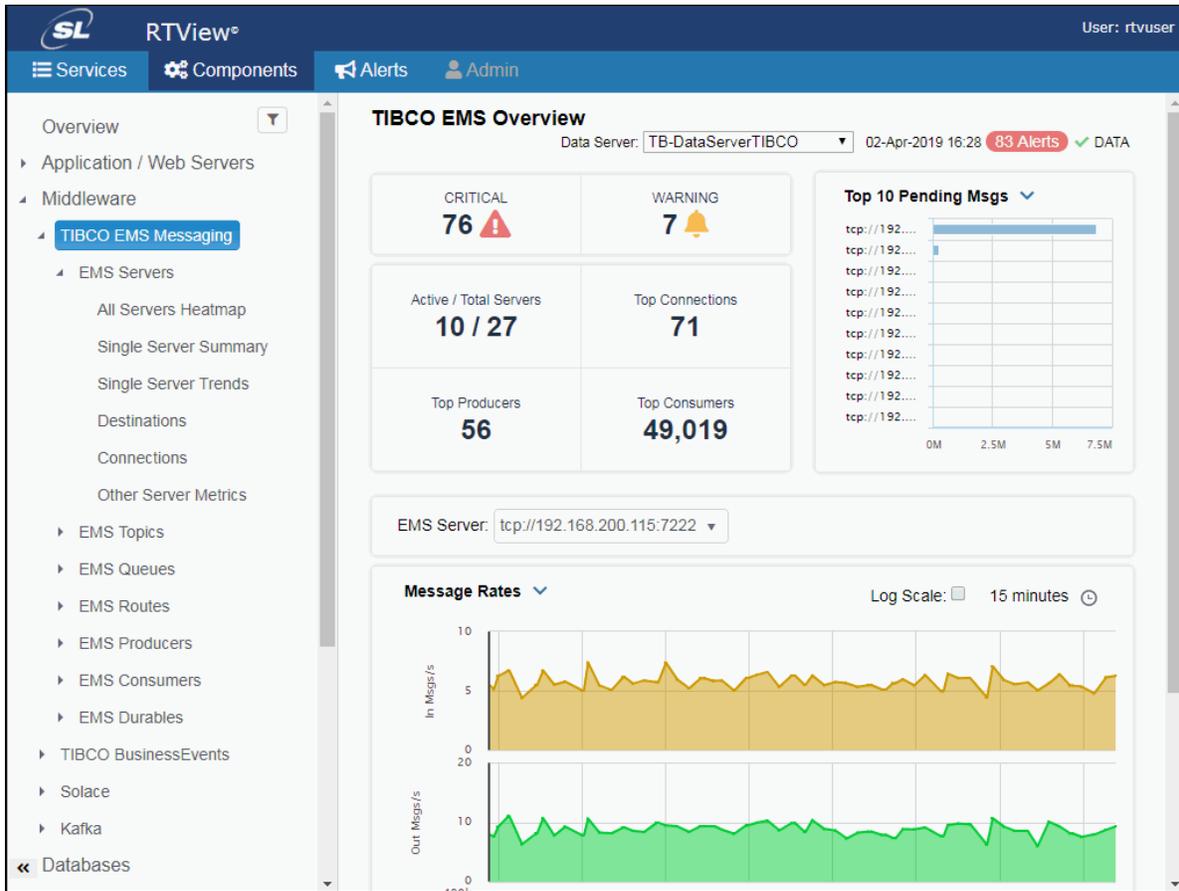
## TIBCO EMS Overview - HTML

The **TIBCO EMS Overview** is the top-level display for the TIBCO Enterprise Message Service Solution Package, which provides a good starting point for immediately getting the status of all your connections on your Data Server. You can select the RTView DataServer for which you want to see data and easily view the current data for that DataServer including:

- The total number of active alerts for the selected DataServer, including the total number of critical and warning alerts.
- The number of active servers and the total number of servers.
- The highest number of connections on a particular server on your connected DataServer.
- The highest number of producers on a particular server on your connected DataServer.
- The highest number of consumers on a particular server on your connected DataServer.
- A visual list of the top 10 servers containing the most total pending messages/connections/incoming messages/Async DB size in bytes on your connected DataServer.
- The total pending messages, the outgoing messages per second, and the incoming messages per second for a selected EMS Server on your connected DataServer.

You can hover over each region in the upper half of the Overview to see more detail. You can also drill down to see even more detail by clicking on each respective region in the Overview. For example, clicking on the alerts in the **CRITICAL** and **WARNING** alerts region opens the **Alerts Table by Components** display.

The bottom half of the display provides a message rates trend graph for a selected EMS server. You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.



## EMS Servers - HTML

These displays present performance metrics and alert status for all EMS servers. Clicking **EMS Servers** from the left/navigation menu opens the ["TIBCO EMS Servers Table - HTML"](#) display, which shows all available utilization metrics for all EMS servers. The options available under **EMS Servers** are:

- **All Servers Heatmap:** Opens the ["TIBCO EMS Servers Heatmap - HTML"](#), which shows server and alert status for all EMS servers.
- **Single Server Summary:** Opens the ["TIBCO EMS Server Summary - HTML"](#) display, which shows information for a single EMS server such as server connection details, the number of client connections, memory utilization, message performance metrics and alert status.
- **Single Server Trends:** Opens the ["TIBCO EMS Server Trends - HTML"](#) display, which shows utilization metrics for a single EMS server, such as the number of client connections, number of pending messages and in/out rate, and memory and disk utilization.
- **Destinations:** Opens the ["TIBCO EMS Server Destinations - HTML"](#) display, which shows destination data for a selected server.
- **Connections:** Opens the ["TIBCO EMS Server Connections - HTML"](#) display, which shows connection information for a selected server.
- **Other Server Metrics:** Opens the ["TIBCO EMS Bridges, Users, Ports - HTML"](#) display, which shows bridges data, along with associated users and ports, for a selected server.

## TIBCO EMS Servers Table - HTML

Investigate detailed utilization metrics for all EMS servers. The **TIBCO EMS Servers Table** contains all metrics available for servers, including the number of current client connections. Each row in the table contains data for a particular server. Click a column header to sort column data in ascending or descending order. Double-click on a table row to drill-down to the ["TIBCO EMS Server Summary - HTML"](#) display and view metrics for that particular server. Toggle between the commonly accessed **Table** and **Heatmap** displays by clicking the drop down list on the display title.

**TIBCO EMS Servers Table** 04-Apr-2019 09:22 No Alerts  DATA

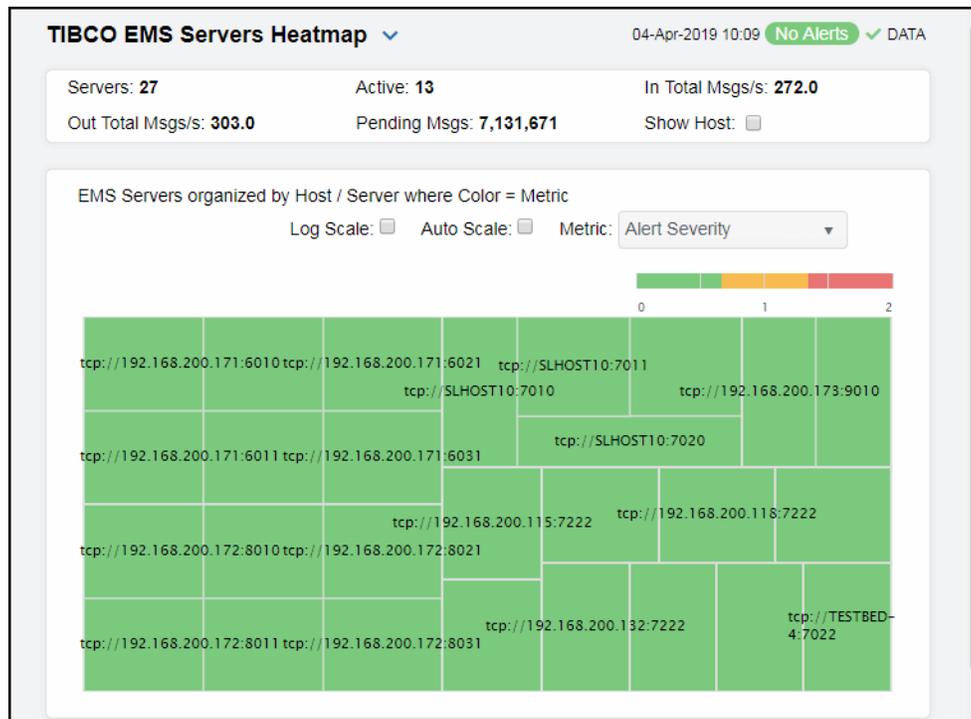
Servers: **27**      Active: **13**      Max Msgs In/s: **16.0**  
 Max Msgs Out/s: **61.0**      Total Pending Msgs: **7,131,652**      Active Only:

URL	Server Name	Host	Expired
tcp://192.168.200.115:7222	EMS-SERVER	192.168.200.115	
tcp://192.168.200.116:7222	EMS-SERVER	192.168.200.116	
tcp://192.168.200.117:7222	EMS-SERVER	192.168.200.117	
tcp://192.168.200.118:7222	EMS-SERVER	192.168.200.118	
tcp://192.168.200.121:7222	EMS-SERVER	192.168.200.121	
tcp://192.168.200.132:7222	Unknown (tcp://192.168.200.132:7222)	192.168.200.132	<span style="color: red;">🔄</span>
tcp://192.168.200.138:7222	Unknown (tcp://192.168.200.138:7222)	192.168.200.138	<span style="color: red;">🔄</span>
tcp://192.168.200.34:7222	EMS-SERVER-TB34	192.168.200.34	
tcp://SLHOST10:7010	Unknown (tcp://SLHOST10:7010)	SLHOST10	<span style="color: red;">🔄</span>
tcp://SLHOST10:7011	Unknown (tcp://SLHOST10:7011)	SLHOST10	<span style="color: red;">🔄</span>
tcp://SLHOST10:7021	Unknown (tcp://SLHOST10:7021)	SLHOST10	<span style="color: red;">🔄</span>
tcp://SLHOST10:7020	Unknown (tcp://SLHOST10:7020)	SLHOST10	<span style="color: red;">🔄</span>
tcp://192.168.200.171:6010	EMS-SLDEMOS2-6010	192.168.200.171	
tcp://192.168.200.171:6011	EMS-SLDEMOS2-6010	192.168.200.171	
tcp://192.168.200.171:6020	EMS-SLDEMOS2-6020	192.168.200.171	

## TIBCO EMS Servers Heatmap - HTML

Clicking **All Servers Heatmap** in the left/navigation menu opens the **TIBCO EMS Servers Heatmap**, which allows you to view the status and alerts of all EMS servers. Use the **Metric** drop-down menu to view the **Alert Severity, Alert Count, Connections, Pending Messages, Inbound Message Rate, Outbound Message Rate, or Message Memory Percent (%)**.

The heatmap is organized by host, each rectangle representing a server. The rectangle color indicates the most critical alert state. Click on a node to drill-down to the "[TIBCO EMS Server Summary - HTML](#)" display and view metrics for a particular server. Toggle between the commonly accessed **Table** and **Heatmap** displays by clicking the drop down list on the display title. Mouse-over rectangles to view more details about host performance and status.



## Available Metrics

Select the metric driving the heatmap display. The default is **Alert Severity**. Each **Metric** has a color gradient bar that maps values to colors. The heatmap organizes the servers by host, where each rectangle represents a server. Mouse-over any rectangle to display the current values of the metrics for the Server. Click on a rectangle to drill-down to the associated ["TIBCO EMS Server Summary - HTML"](#) display for a detailed view of metrics for that particular server.

**Alert Severity** The maximum alert level in the item (index) associated with the rectangle. Values range from **0** to **2**, as indicated in the color gradient bar , where **2** is the greatest **Alert Severity**.

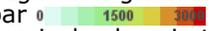
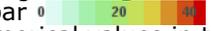
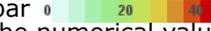
**2** -- Metrics that have exceeded their specified **ALARMLEVEL** threshold and have an Alert Severity value of **2** are shown in red. For a given rectangle, this indicates that one or more metrics have exceeded their alarm threshold.

**1** -- Metrics that have exceeded their specified **WARNINGLEVEL** threshold and have an Alert Severity value of **1** are shown in yellow. For a given rectangle, this indicates that one or more metrics have exceeded their warning threshold.

**0** -- Metrics that have not exceeded either specified threshold have an Alert Severity value of **0** and are shown in green. For a given rectangle, this indicates that no metrics have exceeded a specified alert threshold.

**Alert Count** The total number of alarm and warning alerts in a given item (index) associated with the rectangle.

The color gradient bar shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

- Connections** The total number of connections in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of connections in the heatmap. The middle value in the gradient bar indicates the middle value of the range.  
The **Auto** option does not impact this metric.
- Pending Msgs** The total number of pending messages in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of **EmsServerPendingMsgsHigh**, which is **3500**. The middle value in the gradient bar indicates the middle value of the range (the default is **1750**).  
When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.
- In Msg Rate** The total number of inbound messages in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of **EmsServerInMsgRateHigh**, which is **40**. The middle value in the gradient bar indicates the middle value of the range (the default is **20**).  
When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.
- Out Msg Rate** The total number of outbound messages in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of **EmsServerOutMsgRateHigh**, which is **40**. The middle value in the gradient bar indicates the middle value of the range (the default is **20**).  
When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.
- Mem Msg %** The percent (%) memory used by messages in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of **EmsServerMemUsedHigh**, which is **40**. The middle value in the gradient bar indicates the middle value of the range (the default is **20**).  
When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

## TIBCO EMS Server Summary - HTML

Clicking **Single Server Summary** in the left/navigation menu opens the **TIBCO EMS Server Summary** display, which allows you to track utilization and performance metrics for specific servers. Clicking on the message/connection information boxes at the top of the display takes you to the "[TIBCO EMS Server Destinations - HTML](#)" display, where you can view additional destination data. In the trend graph region, you can select from **Message Rates**, which traces inbound/outbound messages per second, or **Message Flows**, which traces total inbound/outbound messages in bytes. Clicking the **Critical/Warning** link at the bottom of the display opens the Alerts Table by Component display. Clicking the **Topics** link at the bottom of the display opens the "[TIBCO EMS Topics Table - HTML](#)" display. Clicking the **Queues** link at the bottom of the display opens the TIBCO EMS Queues Table display.

### TIBCO EMS Server Summary 04-Apr-2019 10:31 No Alerts DATA

EMS Server: tcp://192.168.200.115:7222

Pending Messages <b>250 K</b>	In Msgs/s <b>7.0</b>	Out Msgs/s <b>9.0</b>	Connections <b>16</b>
Inbound / Outbound KB/s <b>5.2 / 277.7</b>		Message Memory % <b>8.3</b>	

#### Message Rates Log Scale: 15 minutes ⌚

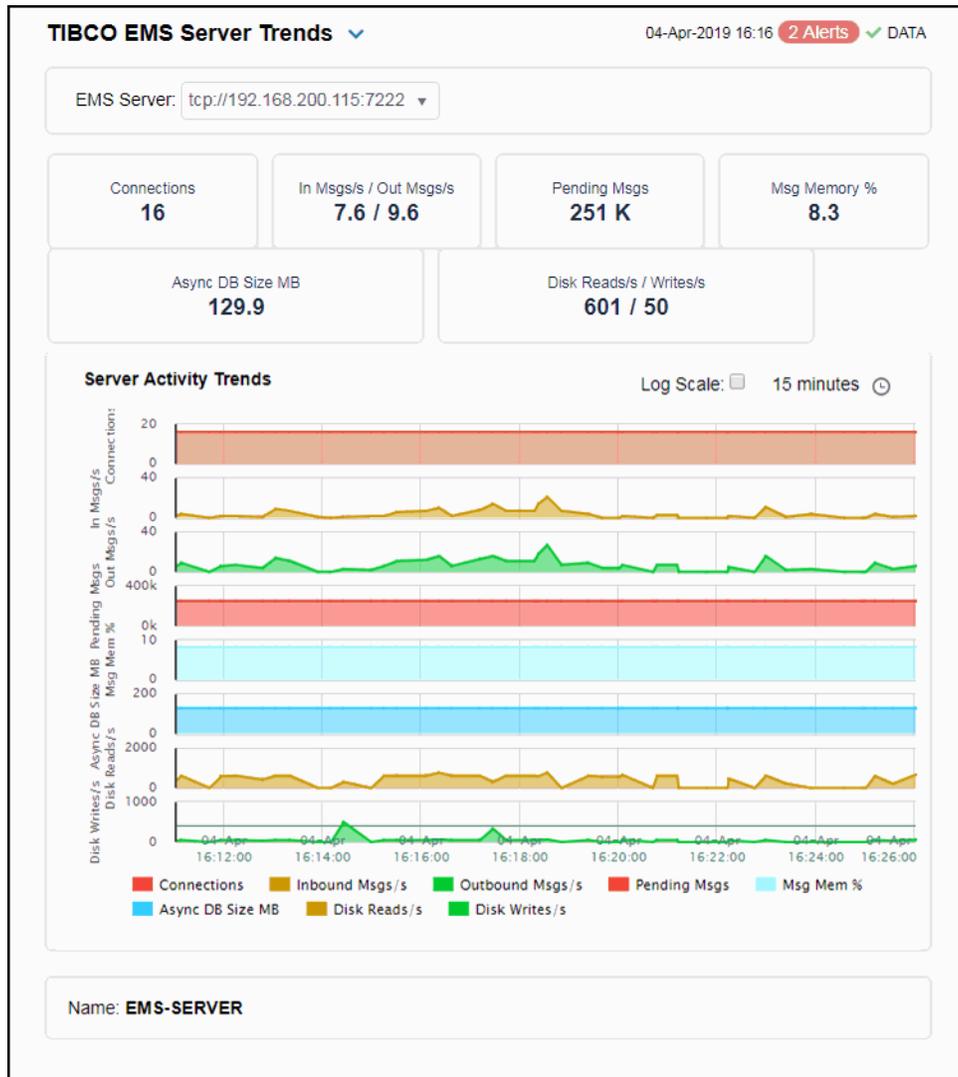
04-Apr 10:22:00 04-Apr 10:24:00 04-Apr 10:26:00 04-Apr 10:28:00 04-Apr 10:30:00 04-Apr 10:32:00 04-Apr 10:34:00

■ In Msgs/s   
 ■ Out Msgs/s   
 ■ Pending Msgs

State: <b>Active</b>	<a href="#">Critical/Warning: 0/0</a>	Backup Name:
Server Name: <b>EMS-SERVER</b>	FT URL:	
<a href="#">Topics: 48,764</a>	<a href="#">Queues: 188</a>	Durables: <b>253</b>
Producers: <b>49</b>	Consumers: <b>49,020</b>	Routes:
Sync DB Size MB: <b>0.0</b>	Start Time: <b>48 years</b>	Uptime: <b>48 years</b>
Async DB Size MB: <b>129.85</b>	Version: <b>6.0.0.8</b>	Message Memory MB: <b>169.44</b>
Max Message Memory MB: <b>2,048.0</b>	Pending Message Size MB: <b>100.49</b>	
Last Update: <b>04-Apr-2019 10:32:44</b>		

## TIBCO EMS Server Trends - HTML

Clicking **Single Server Trends** in the left/navigation menu opens the **TIBCO EMS Server Trends** display, which allows you to view trend graphs in parallel to investigate performance issues for a specific server. Clicking on the message/connection information boxes at the top of the display takes you to the ["TIBCO EMS Servers Table - HTML"](#) display, where you can view additional data for all of the servers. Hovering over the trend graphs displays data for each of the metrics at a specific time.



## TIBCO EMS Server Destinations - HTML

Clicking **Destinations** in the left/navigation menu opens the **TIBCO EMS Server Destinations** display, which allows you to view queue and topic information related to a particular EMS server.

**TIBCO EMS Server Destinations** ▾
16-May-2019 10:51 No Alerts ✓ DATA

EMS Server:

Pending Messages  
**4**

In Msgs/s  
**17.9**

Out Msgs/s  
**12.2**

Connections  
**42**

Inbound / Outbound KB/s  
**0.0 / 0.0**

Message Memory %  
**0.0**

Queue Name	In Msgs/s	In Total Msgs	Out Msgs/s	Out Total
amx.governance.internal.stats	0.0	1,692,765	0.0	1
amx.governance.stats	0.0	310,518	0.0	
cl_logservice_queue	0.0	0	0.0	
cl_payload_queue	0.0	0	0.0	
com.tibco.amf.admin.deploymentServerQu	0.0	0	0.0	
com.tibco.amf.admin.deploymentServerQu	0.0	0	0.0	

Topic Name	In Msgs/s	In Total Msgs	Out Msgs/s	Out Total
adb.custom.jmsrequest	0.0	0	0.0	
adb.standard.jmsrequest	0.0	0	0.0	
rtv.amx.governance.internal.stats	0.0	1,692,765	0.0	
rtv.amx.governance.stats	0.0	310,518	0.0	
sample	0.0	0	0.0	
topic.sample	0.0	0	0.0	

Server Name: **EMS-SERVER**      Last Update: **16-May-2019 10:51:58**

## TIBCO EMS Server Connections - HTML

Clicking **Connections** in the left/navigation menu opens the **TIBCO EMS Server Connections** display, which allows you to view metrics for all connections on a single server. The table lists all of the connections and their associated metrics for the selected server. The bottom portion of the display lists additional details for the selected server. Clicking the **Critical/Warning** link at the bottom of the display opens the Alerts Table by Component display. Clicking the **Topics** link at the bottom of the display opens the “[TIBCO EMS Topics Table - HTML](#)” display. Clicking the **Queues** link at the bottom of the display opens the TIBCO EMS Queues Table display.

TIBCO EMS Server Connections 17-May-2019 16:05 1 Alert DATA

EMS Server: tcp://192.168.200.115:7222

Conn ID	Client ID	connectionURL	User Name	Host
1249	BW-null-topic-ADB_Operations_Ext	[anonymous@S	anonymous	SLHOST15
1422	BW-null-topic-ADB_Operations-AD	[anonymous@S	anonymous	SLHOST15

State: Active [Critical/Warning: 1/0](#) Backup Name:

Server Name: EMS-SERVER FT URL:

[Topics: 48,764](#) [Queues: 185](#) Durables: 253

Producers: 46 Consumers: 49,017 Routes:

Last Update: 17-May-2019 16:14:45

## TIBCO EMS Bridges, Users, Ports - HTML

Clicking **Other Server Metrics** from the left/navigation menu opens the **TIBCO EMS Bridges, Users, Ports** display, which allows you to view bridges configured on an EMS Server, as well as their associated users and ports.

**TIBCO EMS Bridges, Users, Ports** ▾
08-Apr-2019 10:09 2 Alerts ✓ DATA

EMS Server:

**Bridges**

source ▾	target	Se
amx.governance.stats	rtv.amx.governance.stats	
amx.governance.internal.stats	rtv.amx.governance.internal.stats	

**Users**

Description	External	Name
Route Server		EMS-SERVER2
Main Server		EMS-SERVER
Administrator		admin

**Listen Ports**

port	URL
tcp://7222	tcp://192.168.200.115:7222

Name: **EMS-SERVER**

Connections: **2**

## EMS Topics - HTML

These displays present several views of performance metrics for topics. Clicking **EMS Topics** from the left/navigation menu opens the ["TIBCO EMS Topics Table - HTML"](#) display, which shows performance and utilization metrics and trends for all topics defined on a specified server, including consumer and subscriber count, memory utilization, and message performance metrics. You can also view all servers that have a specific topic defined in the ["TIBCO EMS Topic Summary - HTML"](#) display, and you can see a list of all the servers on which those topics are defined on the ["TIBCO EMS Topic Detail by Server - HTML"](#) display. The options available under **EMS Topics** are:

- **All Topics Summary:** Clicking **All Topics Summary** opens the ["TIBCO EMS Topics for Server Summary - HTML"](#) display, which shows performance and utilization metrics and trends for all topics defined on a specified server, including consumer and subscriber count, memory utilization, and message performance metrics.
- **All Topics Heatmap:** Clicking **All Topics Heatmap** opens the ["TIBCO EMS Topics Heatmap - HTML"](#), which is a heatmap representation of a selected set of metrics from Topics organized by Server that allows you to track performance and utilization metrics and trends for all topics on a single server.
- **Single Topic Summary:** Clicking **Single Topic Summary** opens the ["TIBCO EMS Topic Summary - HTML"](#) which shows detailed performance and utilization metrics and trends for a specified topic on a single server, including producer and consumer counts, and message performance metrics.
- **Topic Detail by Server:** Clicking **Topic Detail by Server** opens the ["TIBCO EMS Topic Detail by Server - HTML"](#), which shows performance and utilization metrics for all servers that have a specified topic defined, including consumer and subscriber count, and message performance metrics.

## TIBCO EMS Topics Table - HTML

Clicking **EMS Topics** from the left/navigation menu opens the **TIBCO EMS Topics Table** display, which allows you to track performance and utilization metrics for all topics on a single server. You can enter a string in the **Filter by Topic Name** field to show only topics in the table with names that contain the string. For example, if you enter the string Madrid, all topics with Madrid in the topic name are shown in the table. If no entry is made, all topic names are shown. For most use cases, you can enter a portion of the topic name. Double-clicking on a row in the table opens the data for the selected topic in the "[TIBCO EMS Topic Summary - HTML](#)" display so that you can view additional metrics for the selected topic.

**TIBCO EMS Topics Table** ▾
08-Apr-2019 15:28 2 Alerts ✓ DATA

EMS Server: - All - ▾

Filter by Topic Name:  Shown: **306**  
 Total: **49,454**

Topic Name	URL	Expired	Alert Level	Alert Count	
PROD3.emsmgr.topic.market.cash.flow.tra	tcp://192.168.200.172:8031		✓	0	
PROD3.emsmgr.topic.market.cash.flow.tra	tcp://192.168.200.172:8031		✓	0	
PROD3.emsmgr.topic.market.cash.flow.tra	tcp://192.168.200.172:8031		✓	0	
PROD3.emsmgr.topic.market.cash.flow.tra	tcp://192.168.200.172:8031		✓	0	
PROD3.emsmgr.topic.market.cash.flow.tra	tcp://192.168.200.172:8031		✓	0	
PROD3.emsmgr.topic.market.cash.flow.tra	tcp://192.168.200.172:8031		✓	0	
PROD3.emsmgr.topic.market.cash.flow.tra	tcp://192.168.200.172:8031		✓	0	
PROD3.emsmgr.topic.market.cash.flow.tra	tcp://192.168.200.172:8031		✓	0	
PROD3.emsmgr.topic.market.cash.flow.tra	tcp://192.168.200.172:8031		✓	0	
PROD3.emsmgr.topic.market.cash.flow.tra	tcp://192.168.200.172:8031		✓	0	

⏪ ⏩

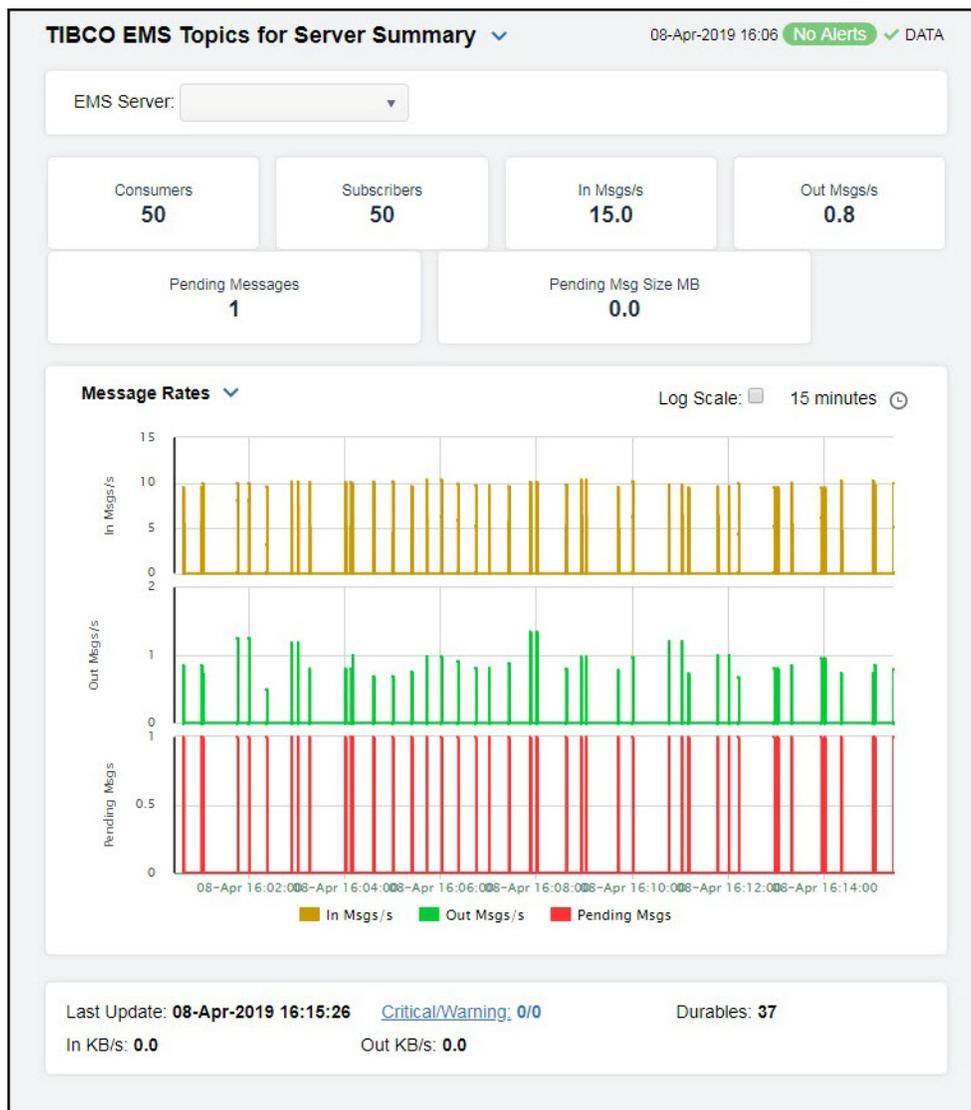
Page  of 8

⏪ ⏩

1 - 40 of 306 items

## TIBCO EMS Topics for Server Summary - HTML

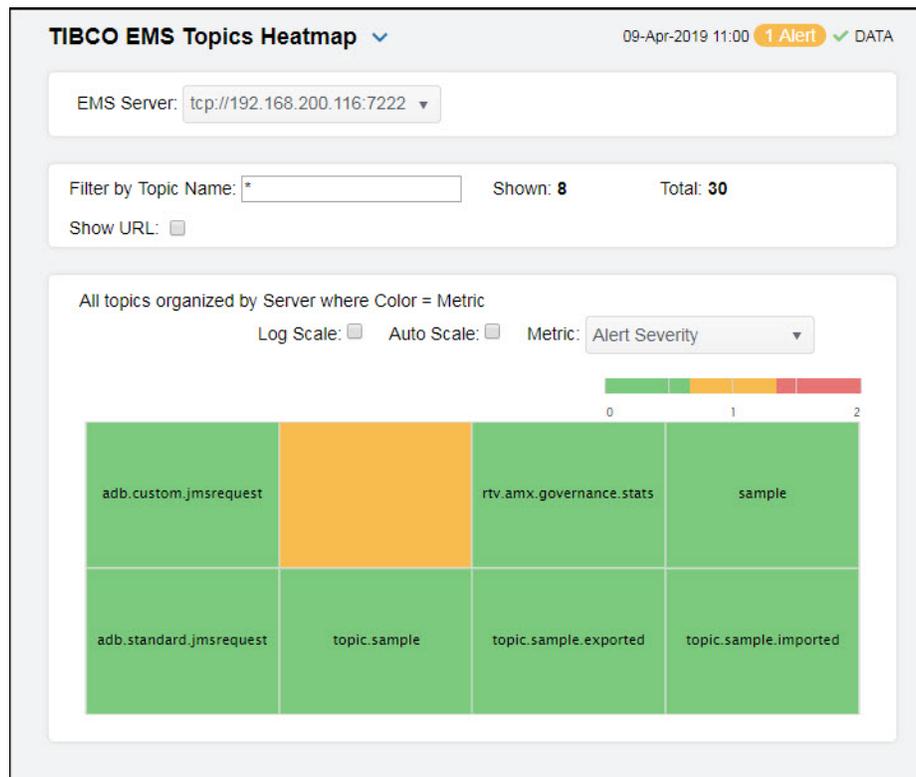
Clicking **All Topics Summary** from the left/navigation menu opens the **TIBCO EMS Topics for Server Summary** display, which allows you to track performance and utilization metrics and trends for all topics on a single server. Clicking on the server information boxes at the top of the display takes you to the ["TIBCO EMS Topics Table - HTML"](#) display, where you can view additional data on all topics. In the trend graph region, you can select from **Message Rates**, which traces inbound/outbound messages per second, **KB Rates**, which traces total inbound/outbound messages per second in kilobytes, or **Pending Msgs**, which traces the total number of messages for all topics on the server currently waiting to be processed and the total size of messages, in megabytes, for all topics on the server currently waiting to be processed. Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display.



## TIBCO EMS Topics Heatmap - HTML

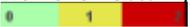
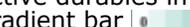
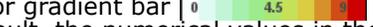
Clicking **All Topics Heatmap** from the left/navigation menu opens the **TIBCO EMS Topics Heatmap**, which is a heatmap representation of a selected set of metrics from Topics organized by Server that allows you to track performance and utilization metrics and trends for all topics on a single server. This heatmap allows you to view status and alerts of all topics for a server. You can enter a string in the **Filter by Topic Name** field to show only topics in the table with names that contain the string. For example, if you enter the string Madrid, all topics with Madrid in the topic name are shown in the table. If no entry is made, all topic names are shown. For most use cases, you can enter a portion of the topic name. Use the **Metric** drop-down menu to view to **Alert Severity, Alert Count, Consumers, Durables, Subscribers, Pending Messages, Inbound Message Rate, Inbound Total Messages, Outbound Message Rate, or Outbound Total Messages**.

The heatmap is organized so that each rectangle represents a Topic on the selected Server. The rectangle color indicates the value of the selected metric in the **Metric** drop down list. You can mouse-over rectangles to view more details about the performance and status of each topic or click on a rectangle to drill-down to the ["TIBCO EMS Topic Summary - HTML"](#) display and view metrics for that particular Topic.



### Available Metrics

Select the metric driving the heatmap display. The default is Alert Severity. Each **Metric** has a color gradient bar that maps values to colors. The heatmap organizes the topics by server, where each rectangle represents a Topic. Mouse-over any rectangle to display the current values of the metrics for the Topic. Click on a rectangle to drill-down to the associated ["TIBCO EMS Topic Summary - HTML"](#) display for a detailed view of metrics for that particular topic.

<b>Alert Severity</b>	<p>The maximum alert level in the item (index) associated with the rectangle. Values range from <b>0</b> to <b>2</b>, as indicated in the color gradient bar , where <b>2</b> is the greatest <b>Alert Severity</b>.</p> <p><b>2</b> -- Metrics that have exceeded their specified <b>ALARMLEVEL</b> threshold and have an Alert Severity value of <b>2</b> are shown in red. For a given rectangle, this indicates that one or more metrics have exceeded their alarm threshold.</p> <p><b>1</b> -- Metrics that have exceeded their specified <b>WARNINGLEVEL</b> threshold and have an Alert Severity value of <b>1</b> are shown in yellow. For a given rectangle, this indicates that one or more metrics have exceeded their warning threshold.</p> <p><b>0</b> -- Metrics that have not exceeded either specified threshold have an Alert Severity value of <b>0</b> and are shown in green. For a given rectangle, this indicates that no metrics have exceeded a specified alert threshold.</p>
<b>Alert Count</b>	<p>The total number of alarm and warning alerts in a given item (index) associated with the rectangle.</p> <p>The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>Consumers</b>	<p>The total number of consumers in a given item (index) associated with the rectangle. The color gradient bar  shows the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of consumers in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p> <p>The <b>Auto Scale</b> option does not impact this metric.</p>
<b>Durables</b>	<p>The total number of active durables in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of durables in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>Subscribers</b>	<p>The total number of subscribers in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of subscribers in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>Pending Msgs</b>	<p>The total number of pending messages in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from <b>0</b> to the alert threshold of <b>EmsTopicsPendingMsgsHigh</b>, which is <b>3000</b>. The middle value in the gradient bar indicates the middle value of the range (the default is <b>1500</b>).</p> <p>When <b>Auto Scale</b> is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>
<b>In Msg /sec</b>	<p>The number of inbound messages per second in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from <b>0</b> to the alert threshold of <b>EmsTopicsInMsgRateHigh</b>, which is <b>9</b>. The middle value in the gradient bar indicates the middle value of the range (the default is <b>4.5</b>).</p> <p>When <b>Auto Scale</b> is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p> <p><b>Note:</b> This metric comes directly from the <b>tibjms.admin.DestinationInfo</b> class from TIBCO.</p>
<b>In Total Msg</b>	<p>The total number of inbound messages in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of receivers in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p> <p>The <b>Auto Scale</b> option does not impact this metric.</p>

**Out Msg/  
sec**

The number of outbound messages per second in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of **EmsTopicsOutMsgRateHigh**, which is **9**. The middle value in the gradient bar indicates the middle value of the range (the default is **4.5**).

When **Auto Scale** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

**Note:** This metric comes directly from the **tibjms.admin.DestinationInfo** class from TIBCO.

**Out Total  
Msgs**

The total number of outbound messages in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of receivers in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

The **Auto Scale** option does not impact this metric.

## TIBCO EMS Topic Summary - HTML

Clicking **Single Topic Summary** from the left/navigation menu opens the **TIBCO EMS Topic Summary** display, which allows you to track performance and utilization metrics for a single topic on a single server. Clicking any of the messages boxes at the top of the display takes you to the ["TIBCO EMS Topics Table - HTML"](#) display, where you can view additional data on all topics. In the trend graph region, you can select from **Message Rates**, which traces inbound/outbound messages per second, or **Message Flows**, which traces total inbound/outbound messages in bytes. Clicking the **Critical/Warning** link at the bottom of the display opens the Alerts Table by Component display.



## TIBCO EMS Topic Detail by Server - HTML

Clicking **Topic Detail by Server** in the left/navigation menu opens the **TIBCO EMS Topic Detail by Server** display, which allows you to track performance and utilization metrics of a single topic across all servers that have the topic defined on it and compare topic activity among servers. Double-clicking any of the rows in the table takes you to the ["TIBCO EMS Topic Summary - HTML"](#) display, where you can view additional data for that particular topic on that particular server.

**TIBCO EMS Topic Detail by Server** 16-May-2019 10:56 Alerts DATA

Topic:

Topics: 20

URL	Expired	In Msgs/s	In Total Msgs	Out Msgs/s	Out Tot
tcp://192.168.200.131:7222		0.0	0	0.0	
tcp://192.168.200.172:8011		0.0	0	0.0	
tcp://192.168.200.172:8010		0.0	0	0.0	
tcp://192.168.200.153:7222		0.0	0	0.0	
tcp://192.168.200.116:7222		0.0	0	0.0	
tcp://192.168.200.34:7222		0.0	0	0.0	
tcp://192.168.200.173:9011		0.0	0	0.0	
tcp://192.168.200.173:9010		0.0	0	0.0	
tcp://192.168.200.118:7222		0.0	0	0.0	
tcp://192.168.200.121:7222		0.0	0	0.0	
tcp://192.168.200.119:7222		0.0	0	0.0	
tcp://192.168.200.171:6011		0.0	0	0.0	
tcp://192.168.200.171:6010		0.0	0	0.0	

## EMS Queues - HTML

These displays present several views of performance metrics for queues. Clicking **EMS Queues** from the left/navigation menu opens the ["TIBCO EMS Queues Table - HTML"](#) display, which shows performance and utilization metrics for all queues defined on a specified server. The options available under **EMS Queues** are:

- **All Queues Summary:** Opens the ["TIBCO EMS Queues for Server Summary - HTML"](#), which shows performance and utilization metrics and trends for all queues defined on a specified server, including message performance metrics.
- **All Queues Heatmap:** Opens the ["TIBCO EMS Queues Heatmap - HTML"](#), which is a heatmap representation of a selected set of metrics that shows performance and utilization metrics and trends for all queues defined on a specified server, including message performance metrics.
- **Single Queue Summary:** Opens the ["TIBCO EMS Queue Summary - HTML"](#), which shows detailed performance and utilization metrics and trends for a specified queue on a single server, including producer and consumer counts, and message performance metrics.
- **Queue Detail by Server:** Opens the ["TIBCO EMS Queue Detail By Server - HTML"](#), which shows performance and utilization metrics for all servers that have a specified queue defined, including consumer and receiver count, as well as message performance metrics.

## TIBCO EMS Queues Table - HTML

Clicking **EMS Queues** from the left/navigation menu opens the **TIBCO EMS Queues Table** display, which allows you to track performance and utilization metrics for all queues on a single server. You can enter a string in the **Filter by Topic Name** field to show only queues in the table with names that contain the string. For example, if you enter the string Madrid, all queues with Madrid in the queue name are shown in the table. If no entry is made, all queue names are shown. For most use cases, you can enter a portion of the queue name. Double-clicking on a row in the table opens the data for the selected queue in the ["TIBCO EMS Queue Summary - HTML"](#) display so that you can view additional metrics for the selected queue.

**TIBCO EMS Queues Table** ▾
10-Apr-2019 16:15 1 Alert ✓ DATA

EMS Server:

Filter by Queue Name:  Shown: **112**  
Total: **187**

Queue Name	URL	Alert Level	Alert Count	In Msgs/s
com.tibco.amr.admin.deploymentServerQu	tcp://192.168.200.115:7222	✓	0	
com.tibco.amf.admin.deploymentServerQu	tcp://192.168.200.115:7222	✓	0	
com.tibco.amf.admin.deploymentServerQu	tcp://192.168.200.115:7222	✓	0	
com.tibco.amf.admin.deploymentServerQu	tcp://192.168.200.115:7222	✓	0	
com.tibco.amf.admin.deploymentServerQu	tcp://192.168.200.115:7222	✓	0	
queue.sample	tcp://192.168.200.115:7222	✓	0	
sample	tcp://192.168.200.115:7222	⚠	1	
sample1	tcp://192.168.200.115:7222	✓	0	
tty0.queue.sample	tcp://192.168.200.115:7222	✓	0	
tty1.queue.sample	tcp://192.168.200.115:7222	✓	0	

Page 1 of 3 1 - 40 of 112 items

## TIBCO EMS Queues for Server Summary - HTML

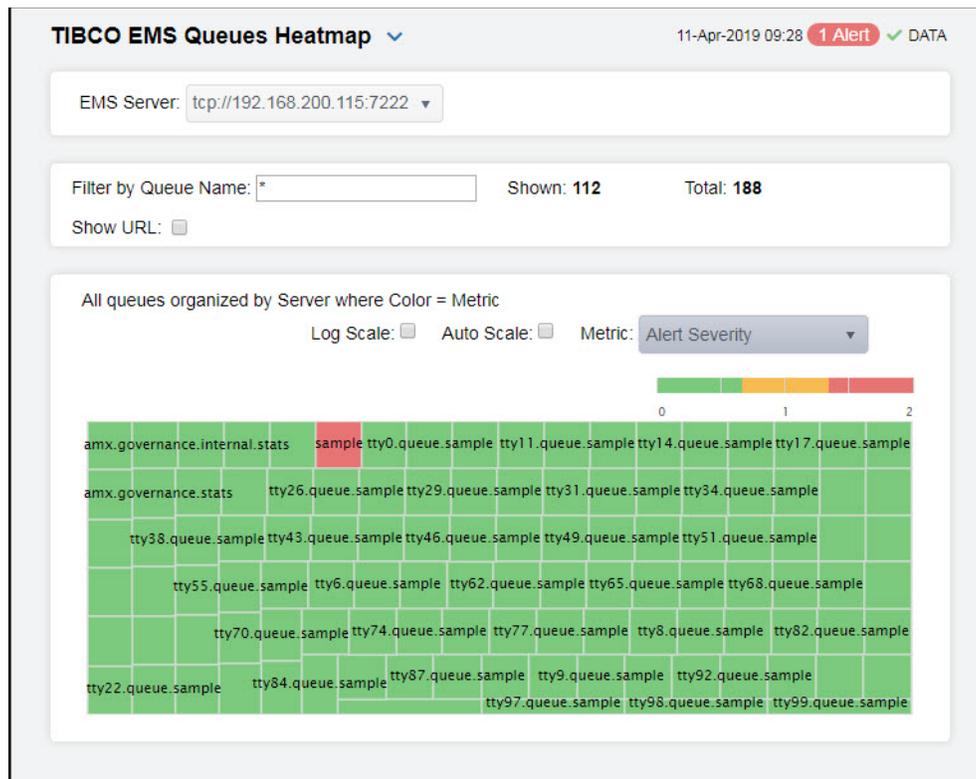
Clicking **All Queues Summary** from the left/navigation menu opens the **TIBCO EMS Queues for Server Summary** displays, which allows you to track performance and utilization metrics and trends for all queues on a single server. Clicking any of the messages boxes at the top of the display takes you to the ["TIBCO EMS Queues Table - HTML"](#) display, where you can view additional data on all queues. In the trend graph region, you can select from **Message Rates**, which traces inbound/outbound messages per second, **KB Rates**, which traces total inbound/outbound messages per second in kilobytes, or **Pending Msgs**, which traces the total number of messages for all queues on the server currently waiting to be processed and the total size of messages, in megabytes, for all queues on the server currently waiting to be processed. Clicking the **Critical/Warning** link at the bottom of the display opens the Alerts Table by Component display.



## TIBCO EMS Queues Heatmap - HTML

Clicking All Queues Heatmap from the left/navigation menu opens the TIBCO EMS Queues Heatmap, which is a heatmap representation of the ["TIBCO EMS Queues Table - HTML"](#) display that allows you to track performance and utilization metrics and trends for all queues on a single server. This heatmap allows you to view status and alerts of all queues for a particular server. You can enter a string in the **Filter by Topic Name** field to show only topics in the table with names that contain the string. For example, if you enter the string Madrid, all topics with Madrid in the topic name are shown in the table. If no entry is made, all topic names are shown. For most use cases, you can enter a portion of the topic name. Use the **Metric** drop-down menu to view to **Alert Severity, Alert Count, Consumers, Receivers, Pending Messages, Inbound Message Rate, Inbound Total Messages, Outbound Message Rate, or Outbound Total Messages**.

The heatmap is organized so that each rectangle represents a queue on the selected server. The rectangle color indicates the most critical alert state. Click on a node to drill-down to the ["TIBCO EMS Queue Summary - HTML"](#) display and view metrics for a particular queue. Toggle between the commonly accessed **Table** (link to the ["TIBCO EMS Queues Table - HTML"](#) display) and **Heatmap** displays. Mouse-over rectangles to view more details about the performance and status of each queue.



## Available Metrics

Select the metric driving the heatmap display. The default is **Alert Severity**. Each Metric has a color gradient bar that maps values to colors. The heatmap organizes the topics by server, where each rectangle represents a Queue. Mouse-over any rectangle to display the current values of the metrics for the Queue. Click on a rectangle to drill-down to the associated ["TIBCO EMS Queue Summary - HTML"](#) display for a detailed view of metrics for that particular queue.

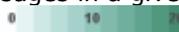
### Alert Severity

The maximum alert level in the item (index) associated with the rectangle. Values range from **0** to **2**, as indicated in the color gradient bar , where **2** is the greatest **Alert Severity**.

-- Metrics that have exceeded their specified **ALARMLEVEL** threshold and have an Alert Severity value of **2** are shown in red. For a given rectangle, this indicates that one or more metrics have exceeded their alarm threshold.

**1** -- Metrics that have exceeded their specified **WARNINGLEVEL** threshold and have an Alert Severity value of **1** are shown in yellow. For a given rectangle, this indicates that one or more metrics have exceeded their warning threshold.

-- Metrics that have not exceeded either specified threshold have an Alert Severity value of **0** and are shown in green. For a given rectangle, this indicates that no metrics have exceeded a specified alert threshold.

<b>Alert Count</b>	<p>The total number of alarm and warning alerts in a given item (index) associated with the rectangle.</p> <p>The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>Consumers</b>	<p>The total number of consumers in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of receivers in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p> <p>The <b>Auto</b> option does not impact this metric.</p>
<b>Receivers</b>	<p>The total number of receivers in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of receivers in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p> <p>The <b>Auto</b> option does not impact this metric.</p>
<b>Pending Msgs</b>	<p>The total number of pending messages in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from <b>0</b> to the alert threshold of <b>EmsQueuesPendingMsgsHigh</b>, which is <b>3000</b>. The middle value in the gradient bar indicates the middle value of the range (the default is <b>1500</b>).</p> <p>When <b>Auto</b> is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p>
<b>In Msgs / sec</b>	<p>The number of inbound messages per second in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from <b>0</b> to the alert threshold of <b>EmsQueuesInMsgRateHigh</b>, which is <b>9</b>. The middle value in the gradient bar indicates the middle value of the range (the default is <b>4.5</b>).</p> <p>When <b>Auto</b> is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p> <p><b>Note:</b> This metric comes directly from the <b>tibjms.admin.DestinationInfo</b> class from TIBCO.</p>
<b>In Total Msg</b>	<p>The total number of inbound messages in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of receivers in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p> <p>The <b>Auto</b> option does not impact this metric.</p>
<b>Out Msgs/ sec</b>	<p>The number of outbound messages per second in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from <b>0</b> to the alert threshold of <b>EmsQueuesOutMsgRateHigh</b>, which is <b>9</b>. The middle value in the gradient bar indicates the middle value of the range (the default is <b>4.5</b>).</p> <p>When <b>Auto</b> is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.</p> <p><b>Note:</b> This metric comes directly from the <b>tibjms.admin.DestinationInfo</b> class from TIBCO.</p>
<b>Out Total Msgs</b>	<p>The total number of outbound messages in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of receivers in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p> <p>The <b>Auto</b> option does not impact this metric.</p>

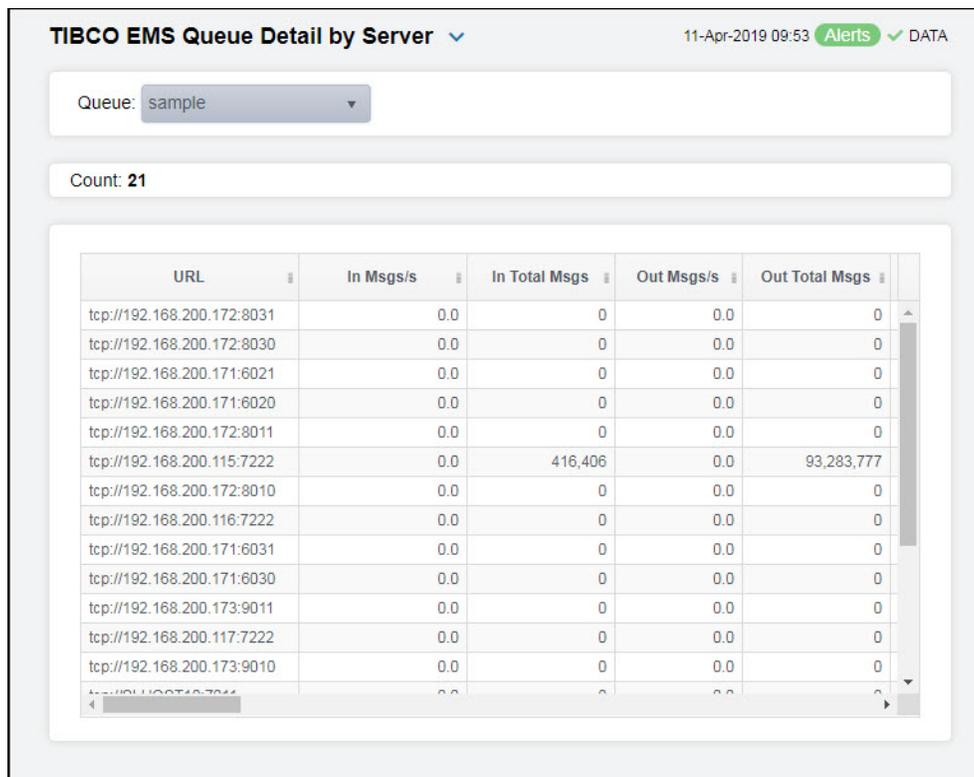
## TIBCO EMS Queue Summary - HTML

Clicking **Single Queue Summary** from the left/navigation menu opens the **TIBCO EMS Queue Summary** display, which allows you to track performance and utilization metrics for a single queue on a single server. Clicking any of the messages boxes at the top of the display takes you to the "[TIBCO EMS Queues Table - HTML](#)" display, where you can view additional data on all topics. In the trend graph region, you can select from **Message Rates**, which traces inbound/outbound messages per second, or **Message Flows**, which traces total inbound/outbound messages in bytes. Clicking the **Critical/Warning** link at the bottom of the display opens the Alerts Table by Component display.



## TIBCO EMS Queue Detail By Server - HTML

Clicking **Queue Detail by Server** in the left/navigation menu opens the **TIBCO EMS Queue Detail by Server** display, which allows you to track performance and utilization metrics of a single queue across all servers and compare queue activity among servers. Double-clicking any of the rows in the table takes you to the ["TIBCO EMS Queue Summary - HTML"](#) display, where you can view additional data for that particular queue on that particular server.



**TIBCO EMS Queue Detail by Server** 11-Apr-2019 09:53 Alerts DATA

Queue:

Count: 21

URL	In Msgs/s	In Total Msgs	Out Msgs/s	Out Total Msgs
tcp://192.168.200.172:8031	0.0	0	0.0	0
tcp://192.168.200.172:8030	0.0	0	0.0	0
tcp://192.168.200.171:6021	0.0	0	0.0	0
tcp://192.168.200.171:6020	0.0	0	0.0	0
tcp://192.168.200.172:8011	0.0	0	0.0	0
tcp://192.168.200.115:7222	0.0	416,406	0.0	93,283,777
tcp://192.168.200.172:8010	0.0	0	0.0	0
tcp://192.168.200.116:7222	0.0	0	0.0	0
tcp://192.168.200.171:6031	0.0	0	0.0	0
tcp://192.168.200.171:6030	0.0	0	0.0	0
tcp://192.168.200.173:9011	0.0	0	0.0	0
tcp://192.168.200.117:7222	0.0	0	0.0	0
tcp://192.168.200.173:9010	0.0	0	0.0	0
tcp://192.168.200.173:9011	0.0	0	0.0	0

## EMS Routes - HTML

These displays present performance metrics and alert status for all routes or one route on an EMS Server. Clicking **EMS Routes** from the left/navigation menu opens the ["TIBCO EMS Routes - HTML"](#) display, which shows all available utilization metrics for all EMS routes on a specific EMS server. The option available under **EMS Routes** is:

- **Route Summary:** Opens the ["TIBCO EMS Route - HTML"](#) display, which shows metrics and trend data for a particular route on a particular EMS Server.

## TIBCO EMS Routes - HTML

Clicking **EMS Routes** from the left/navigation menu opens the **TIBCO EMS Routes** display, which shows all available utilization metrics for all routes on a specific EMS server. Double-clicking on a route in the Routes for Server table opens the ["TIBCO EMS Route - HTML"](#) display, which shows additional details for the selected route. Inbound metrics, such as **In Msgs/s**, indicate an in route to the server. Outbound metrics, such as **Out Msgs/s**, indicate an out route to the server.

**TIBCO EMS Routes** 11-Apr-2019 11:11 No Alerts DATA

EMS Server:  FT Standby

**Totals for Server**  
 In Msgs/s: **0.0**      In Total Msgs: **0**      Out Msgs/s: **0.0**  
 Out Total Msgs: **0**

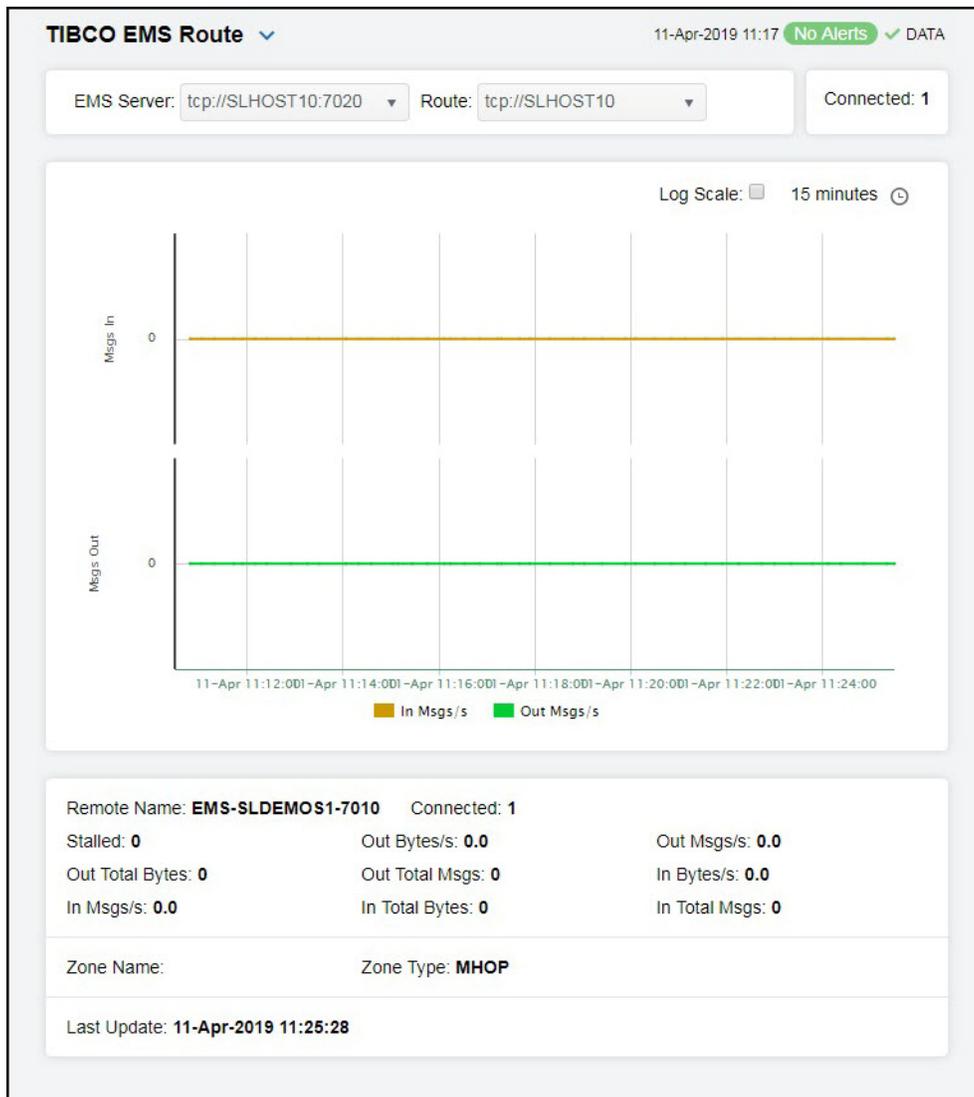
Routes for Server: **EMS-SLDEMOS1-7010**

Remote URL	Remote Name	Connected	Stalled
tcp://SLHOST10:7020,tcp://SLHOST10:7021	EMS-SLDEMOS1-7020	0	
tcp://192.168.200.171:6010,tcp://192.168.200.171:6010	EMS-SLDEMOS2-6010	0	

Last Update: 11-Apr-2019 11:19:53

## TIBCO EMS Route - HTML

Clicking **Route Summary** from the left/navigation menu opens the **TIBCO EMS Route** display, which shows metrics and trend data for a particular route on a particular EMS Server. Hovering over the trend graphs displays data for each of the metrics at a specific time.



## EMS Producers - HTML

These displays present performance metrics and alert status for all producers or one producer on an EMS Server. Clicking **EMS Producers** from the left/navigation menu opens the ["TIBCO EMS Producers - HTML"](#) display, which shows all available utilization metrics for all EMS producers on a specific EMS server. The option available under **EMS Producers** is:

- **Producer Summary:** Opens the ["TIBCO EMS Producer - HTML"](#) display, which shows metrics and trend data for a particular producer on a particular EMS Server.

## TIBCO EMS Producers - HTML

Clicking **EMS Producers** from the left/navigation menu opens the **TIBCO EMS Producers** display, which shows utilization metrics for all producers on a particular EMS Server. You can filter the list of producers in the **Producers for Server** table by **Client ID** and/or **Destination**. Clicking the **Topics** link in the bottom portion of the display opens the "[TIBCO EMS Topics Table - HTML](#)" display. Clicking the **Queues** link in the bottom portion of the display opens the "[TIBCO EMS Queues Table - HTML](#)" display.

**TIBCO EMS Producers** ▾
11-Apr-2019 13:29 No Alerts ✓ DATA

EMS Server: 
Active

Client ID: 
RegEx: 
Destination:

RegEx:

Count: **6**

Msgs/s: **0.0**

Msgs Total: **99,492,972**

Bytes/s: **0.0**

Total Bytes: **236,200,115,112**

Producers for Server: **EMS-SERVER**

ID	Client ID	Destination Name	Msgs/s	Msgs Total
584234		amx.governance.stat	0.0	14,569,268
584258		amx.governance.inter	0.0	84,066,660
584210		amx.governance.stat	0.0	857,044
584224		com.tibco.amf.admin.	0.0	0
584230		com.tibco.amf.admin.	0.0	0
584227		com.tibco.amf.admin.	0.0	0

[Topics: 30](#)  
 Producers: **30**

[Queues: 242](#)  
 Consumers: **88**

Durables: **5**  
 Routes: **1**

Last Update: **11-Apr-2019 13:38:15**

## TIBCO EMS Producer - HTML

Clicking **Producer Summary** from the left/navigation menu opens the **TIBCO EMS Producer** display, which shows metrics and trend data for a particular producer on a particular EMS Server. Hovering over the trend graphs displays data for each of the metrics at a specific time.



## EMS Consumers - HTML

These displays present performance metrics and alert status for all consumers or one consumer on an EMS Server. Clicking **EMS Consumers** from the left/navigation menu opens the "[TIBCO EMS Consumers - HTML](#)" display, which shows all available utilization metrics for all EMS consumers on a specific EMS server. The option available under **EMS Consumers** is:

- **Consumer Summary:** Opens the "[TIBCO EMS Consumer - HTML](#)" display, which shows metrics and trend data for a particular consumer on a particular EMS Server.

## TIBCO EMS Consumers - HTML

Clicking **EMS Consumers** from the left/navigation menu opens the **TIBCO EMS Consumers** display, which shows utilization metrics for all consumers on a particular EMS Server. You can filter the list of consumers in the **Consumers for Server** table by **Client ID** and/or **Destination**. Clicking the **Topics** link in the bottom portion of the display opens the "[TIBCO EMS Topics Table - HTML](#)" display. Clicking the **Queues** link in the bottom portion of the display opens the "[TIBCO EMS Queues Table - HTML](#)" display.

**TIBCO EMS Consumers** ▾
11-Apr-2019 13:55 No Alerts DATA

EMS Server: tcp://192.168.200.116:7222 ▾

Active

Client ID: \*

RegEx:

Destination: \*

RegEx:

Count: **13**

Consumer Msgs/s: **0.0**

Consumer Total Msgs: **57,470,741**

Consumer Bytes/s: **0.0**

Consumer Total Bytes: **133,402,130,095**

Consumers for Server: **EMS-SERVER**

ID	Client ID	Destination Name	Consumer Msgs/s	Consume Total Msgs
193157		rtv.amx.governance.s	0.0	7,714.1
193175		amx.governance.intel	0.0	42,041.1
193176		amx.governance.stat	0.0	7,715.1
193104		cl_payload_queue	0.0	
193120		cl_logservice_queue	0.0	
192368		adb.custom.jmsreque	0.0	

[Topics: 30](#)

[Queues: 242](#)

Durables: **5**

Producers: **30**

Consumers: **88**

Routes: **1**

Last Update: **11-Apr-2019 14:04:35**

## TIBCO EMS Consumer - HTML

Clicking **Consumer Summary** from the left/navigation menu opens the **TIBCO EMS Consumer** display, which shows metrics and trend data for a particular consumer on a particular EMS Server. Hovering over the trend graphs displays data for each of the metrics at a specific time.



## EMS Durables - HTML

These displays present performance metrics and alert status for all durables or one durable on an EMS Server. Clicking **EMS Durables** from the left/navigation menu opens the “[TIBCO EMS Durables - HTML](#)” display, which shows all available utilization metrics for all EMS durables on a specific EMS server. The option available under **EMS Durables** is:

- **Durable Summary:** Opens the “[TIBCO EMS Durable - HTML](#)” display, which shows metrics and trend data for a particular durable on a particular EMS Server.

## TIBCO EMS Durables - HTML

Clicking **EMS Durables** from the left/navigation menu opens the **TIBCO EMS Durables** display, which shows utilization metrics for all durables on a particular EMS Server. Double-clicking a row in the table opens the selected durable in the “[TIBCO EMS Durable - HTML](#)” display.

**TIBCO EMS Durables** ▾
11-Apr-2019 14:19 No Alerts ✓ DATA

EMS Server: tcp://192.168.200.116:7222 ▾

Active

Total Pending Msgs: 0
Total Pending Size: 0

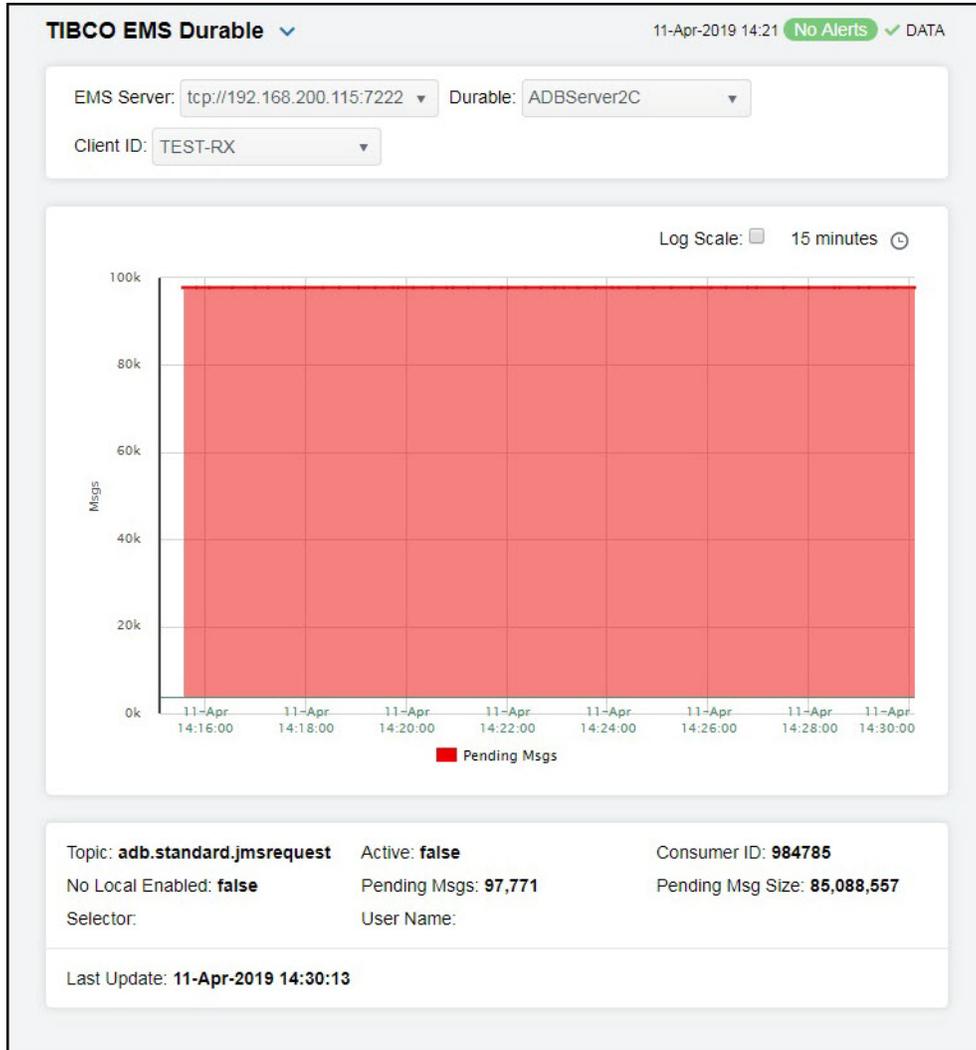
Routes for Server: **EMS-SERVER**

Name	Topic	Active	Client ID	Consumer ID	No I
ADBServer2C	adb.standard.jmsreque		null	192369	
ADBServer2A	adb.standard.jmsreque		null	192371	
ADBServer2B	adb.standard.jmsreque		null	192370	
ADBServer1	adb.custom.jmsreques		null	192368	
ADBServer	adb.standard.jmsreque		null	192372	

Last Update: 11-Apr-2019 14:27:44

### TIBCO EMS Durable - HTML

Clicking **Durable Summary** from the left/navigation menu opens the **TIBCO EMS Durable** display, which shows metrics and trend data for a particular durable on a particular EMS Server. Hovering over the trend graphs displays data for each of the metrics at a specific time.



---

## TIBCO FTL

The Solution Package for TIBCO FTL Views (and their associated displays) can be found under **Components** tab > **Middleware**. The displays within the Views will be populated with data once the Solution Package for TIBCO FTL is configured in the RTView DataServer for TIBCO and the RTView DataServer for TIBCO is connected to RTViewCentral. The available Views are:

- "FTL Servers"
- "FTL Clients"
- "FTL Events"

### FTL Servers

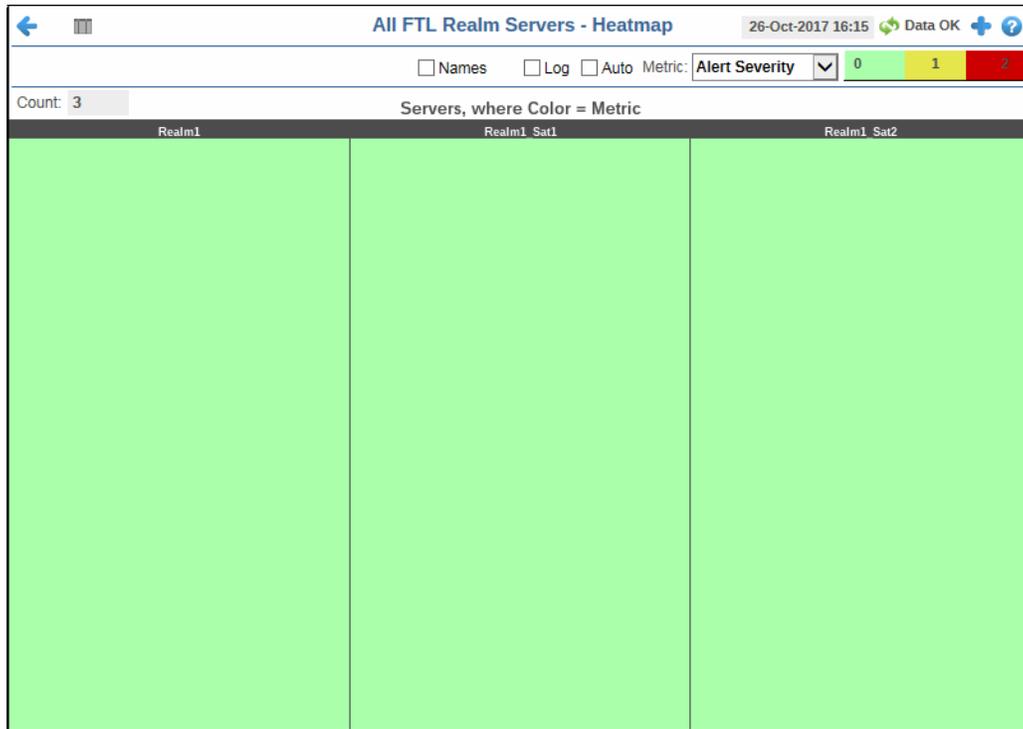
These displays present performance metrics and alert status for all FTL servers. Displays are:

- "All Servers Heatmap": Heatmap shows server and alert status for all FTL servers in all realms.
- "All Servers Table": Table shows all available utilization metrics for all FTL servers.
- "All Group Servers Table": Table shows the status and ID of all FTL Group Servers.
- "All Satellites Table": Table shows the status and ID of all satellites.
- "Server Summary": Current and historical metrics for a single FTL server.

### All Servers Heatmap

This heatmap display provides an easy-to-view interface that allows you to quickly identify the current status of each of your servers for each available metric. You can view the servers in the heatmap based on the following metrics: the current alert severity, the current alert count, the total number of clients, the current amount of CPU being used, the current amount of memory being used for processing, the current amount of virtual memory being used for processing, and the number of inbox send faults. By default, this display shows the heatmap based on the **Alert Severity** metric.

Each heatmap rectangle represents a server. The rectangle color indicates the most critical alert state. You can click on a rectangle to drill-down to the "Server Summary" display and view metrics for that server. Clicking on the  icon in the upper left hand corner of the display toggles between the commonly accessed **Table** and **Heatmap** displays. You can also mouse-over the rectangles to view more details about host performance and status.



#### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

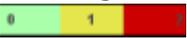
Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

#### Fields and Data

This display includes:

- Names** Select to display the names of servers on the hosts.
- Log** This option enables visualization on a logarithmic scale, and should be used when the range in your data is very broad. For example, if you have data that ranges from the tens to the thousands, then data in the range of tens will be neglected visually if you do not check this option. This option makes data on both extreme ranges visible by using the logarithmic of the values rather than the actual values.
- Auto** When checked, the values of the selected metric are auto-scaled to its highest defined value. When unchecked, the values of the selected metric display based on the threshold defined for the alert associated with the selected metric. Selecting Auto helps to visualize the range of the values currently present for the selected metric instead of the threshold of the alert that has been associated with the metric. All metrics that have not been associated in the heatmap defaults with alerts use a monochromatic color gradient bar (whites and greens). All metrics that have been associated in the heatmap defaults with alerts use a multi-chromatic color gradient bar (reds, yellows, white, and greens).

**Metric** Select the metric driving the heatmap display. The default is Alert Severity. Each **Metric** has a color gradient bar that maps values to colors. The heatmap organizes the servers by host, where each rectangle represents a server. Mouse-over any rectangle to display the current values of the metrics for the Server. Click on a rectangle to drill-down to the associated "Server Summary" display for a detailed view of metrics for that particular server.

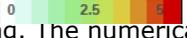
**Alert Severity** The maximum alert level in the item (index) associated with the rectangle. Values range from **0** to **2**, as indicated in the color gradient bar , where **2** is the greatest **Alert Severity**.

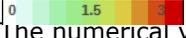
**2** Metrics that have exceeded their specified **ALARMLEVEL** threshold and have an Alert Severity value of **2** are shown in red. For a given rectangle, this indicates that one or more metrics have exceeded their alarm threshold.

**1** Metrics that have exceeded their specified **WARNINGLEVEL** threshold and have an Alert Severity value of **1** are shown in yellow. For a given rectangle, this indicates that one or more metrics have exceeded their warning threshold.

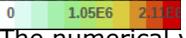
**0** Metrics that have not exceeded either specified threshold have an Alert Severity value of **0** and are shown in green. For a given rectangle, this indicates that no metrics have exceeded a specified alert threshold.

**Alert Count** The total number of alarm and warning alerts in a given item (index) associated with the rectangle.  
The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

**# Clients** The total number of clients in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the defined alert threshold of **TftlServerClientCountHigh**. The middle value in the gradient bar indicates the middle value of the range.

**CPU Usage** The total amount of CPU used. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the defined alert threshold of **TftlServerCpuUsageHigh**. The middle value in the gradient bar indicates the middle value of the range.

**Memory** The current memory being used. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the defined alert threshold of **TftlServerMemoryHigh**. The middle value in the gradient bar indicates the middle value of the range.

**V(irtual) Memory** The current virtual memory being used. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the defined alert threshold of **TftlServerVirtualMemoryHigh**. The middle value in the gradient bar indicates the middle value of the range.

**#Inbox Faults** The total number of inbox faults. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the defined alert threshold of **TftlServerInboxSendFaultsHigh**. The middle value in the gradient bar indicates the middle value of the range.

## All Servers Table

Investigate detailed utilization metrics for all FTL servers. The **All Servers Table** contains all metrics available for servers, including the number of current client connections. Each row in the table describes a different server, and clicking on a table row drills-down to the “[Server Summary](#)” display allowing you to view metrics for that particular server. Clicking the  icon in the upper left hand corner of the display toggles between the commonly accessed **Table** and **Heatmap** displays. You can click a column header to sort column data in numerical or alphabetical order.

Connection	Current Realm Server	Alert Level	Alert Count	On Backup	Backup Server Status	% CPU	Client C
Realm1	192.168.200.227:8050		0	<input type="checkbox"/>	OK	0.70	
Realm1_Sat1	192.168.200.228:8050		0	<input type="checkbox"/>	OK	0.40	
Realm1_Sat2	192.168.200.229:8050		0	<input type="checkbox"/>	OK	2.00	

### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by TIBCO FTL. Refer to the TIBCO FTL documentation for more information regarding these fields.

## Fields and Data

This display includes:

**Count** The total number of active, inactive, and standby FTL servers. Column values are for the server except where noted. **Inactive Servers** are shown in dark red. **Standby Servers** are shown in blue. **Inactive Servers** are shown in dark gray. **Active Backup Servers** are shown in yellow.

### All Servers Table

<b>Connection</b>	The name of the connection.
<b>Current Realm Server</b>	The IP address and port number for the currently connected realm server.
<b>Alert Level</b>	<p>The maximum alert level in the item (index) associated with the rectangle. Values range from 0 to 2, as indicated in the color gradient bar, where <b>2</b> is the greatest Alert Severity.</p> <ul style="list-style-type: none"> <li> One or more alerts have exceeded their specified <b>ALARMLEVEL</b> threshold, have an Alert Severity value of <b>2</b>, and are shown in red.</li> <li> One or more alerts have exceeded their specified <b>WARNINGLEVEL</b> threshold, have an Alert Severity value of <b>1</b>, and are shown in yellow.</li> <li> No alerts have exceeded an alert threshold, which have an Alert Severity value of <b>0</b>, and are shown in green.</li> </ul>
<b>Alert Count</b>	The number of alerts currently on the server.
<b>On Backup</b>	When checked, the server is in backup server mode.*
<b>Backup Server Status</b>	The current backup server status.*
<b>%CPU</b>	The percent CPU used on the server.*
<b>Client Count</b>	<p>The number of clients currently connected to the server.*</p> <p>Note that the client count might not match the number of clients found in the "All Clients Table", possibly due to the following:</p> <ul style="list-style-type: none"> <li>• One client might have one or more group joins resulting in a higher client count. For example, if a client has two group joins, CLIENT_COUNT equals 3, but will only be listed as a single client in the "All Clients Table".</li> <li>• A TIBCO bridge could have one or more logical bridges running inside the bridge process, which could result in an increased CLIENT_COUNT even though there is actually only one client.</li> <li>• Other FTL services could get a Client ID and, hence, be included in the CLIENT_COUNT even though they are not necessarily clients.</li> </ul>
<b>Clients Running</b>	The number of connected clients on the server that currently have a status of RUNNING (which can be less than or equal to the client count).*
<b>Cumulative Client Connects</b>	The total number of clients the server has connected since the server was last started.*

<b>Process Peak RSS (KB)</b>	The maximum RSS memory used, in kilobytes.*
<b>Process RSS (KB)</b>	The current RSS memory used, in kilobytes.*
<b>Process VM (KB)</b>	The current virtual memory used, in kilobytes.*
<b>#Bridge Servers</b>	The number of bridge servers connected.*
<b>#EFTL Clusters</b>	The number of EFTL clusters connected.*
<b>#Group Clients</b>	The number of group clients connected.*
<b>#Group Servers</b>	The number of group servers connected.*
<b>#Persistence Servers</b>	The number of persistence servers connected.*
<b>#Satellites</b>	The number of satellites connected.*
<b>Inbox Send Faults</b>	The total number of faults when sending messages to inbox subscribers.*
<b>Delta Inbox Send Faults</b>	The number of faults when sending messages to inbox subscribers since the last data update.
<b>Rate Inbox Send Faults</b>	The rate of faults when sending messages to inbox subscribers
<b>User CPU Time</b>	Total amount of time the CPU spent, in microseconds, processing object code for users.*
<b>System CPU Time</b>	Total amount of time the CPU spent, in microseconds, processing operating system calls.*
<b>Clients Destroyed</b>	The total number of destroyed clients since the server was last started.*
<b>Client Exceptions</b>	The total number of client exceptions since the server was last started.*
<b>Clients Needing Restart</b>	The total number of clients that had to reconnect since the server was last started.*
<b>Clients Off-line</b>	The number of clients currently offline.*
<b>Clients Out of Sync</b>	The number of clients currently out of sync.*
<b>Client Reconnects</b>	The total number of clients that had to reconnect since the server was last started.*
<b>Client Time-outs</b>	The total number of clients connections that time out due to inactivity since the server was last started.*
<b>Primary Realm Server</b>	The configured primary realm server.*
<b>Backup Realm Server</b>	The configured secondary realm server.*
<b>Server ID</b>	The unique server ID.*
<b>Uptime</b>	The number of days, hours and minutes since the server was last started.*
<b>Version</b>	The FTL version on the server.*
<b>Source</b>	The source of the incoming data.

**Expired**

When checked (Expired=true), monitoring data for the FTL Server row has not been received within the time specified for expiration, which is defined (in seconds) using the **\$ftlServerRowExpirationTime** substitution located in the **conf\rtvapm\_tftlmon.properties** file. If the row has been expired for an extended period of time, the **\$ftlRowExpirationTimeForDelete** substitution determines when the row will be deleted from the cache that drives the display. The default values for the substitutions are 10 and 3600 seconds respectively, meaning that each of the FTL Server rows will have Expired set to true after 10 seconds of inactivity and that expired FTL Server rows will be removed from the cache after 3600 seconds (one hour) of inactivity.

To edit the default/current values, copy the lines below from **rtvapm\_tftlmon.properties** file, paste them into the **sample.properties** file, and modify the lines in the **sample.properties** file:

```
#####
# CACHE / HISTORIAN SETTINGS

sl.rtvview.sub=$ftlServerRowExpirationTime:10
sl.rtvview.sub=$ftlRowExpirationTimeForDelete:0
```

**Data Timestamp**

The date and time this row of data occurred in FTL.\*

**Timestamp**

The date and time this row of data was last updated in RTView.

## All Group Servers Table

This display allows you to view the status of all FTL group servers, see their parent realm, and see which are primary and secondary backup servers. Each row in the table is a different parent realm, and clicking on a table row drills-down to the “Clients by Group” display allowing you to view detailed metrics for that group. You can click a column header to sort column data in ascending or descending order.

Parent Realm	server_mode	server_state	Expired	Timestamp
Realm1	Primary	RUNNING	<input type="checkbox"/>	24-Oct-2017 17:06:02

### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by TIBCO FTL. Refer to the TIBCO FTL documentation for more information regarding these fields.

### Filter By:

**Server:** Select the server for which you want to view data, or select **All Group Servers** to view data for all servers.

**Count** The total number of group servers. Inactive Servers are shown with a dark gray background and Backup Servers are shown with a yellow background.

### Group Server Status Table

**server\_mode** The server mode. The server could be running as a primary server, acting as a backup server or running as a satellite. Valid values are:\*

**Primary** -- the server is running as a primary server.

**Secondary** -- the server is acting as a secondary server.

**server\_state** The current server state. Valid values are:\*

**Running** -- the server is up and running.

**Stopped** -- the server is stopped.

**Expired** When checked (Expired=true), monitoring data for the FTL Server row has not been received within the time specified for expiration, which is defined (in seconds) using the **\$ftlServerRowExpirationTime** substitution located in the **conf\rtvadm\_tftlmon.properties** file. If the row has been expired for an extended period of time, the **\$ftlRowExpirationTimeForDelete** substitution determines when the row will be deleted from the cache that drives the display. The default values for the substitutions are 10 and 3600 seconds respectively, meaning that each of the FTL Server rows will have Expired set to true after 10 seconds of inactivity and that expired FTL Server rows will be removed from the cache after 3600 seconds (one hour) of inactivity.

To edit the default/current values, copy the lines below from **rtvadm\_tftlmon.properties** file, paste them into the **sample.properties** file, and modify the lines in the **sample.properties** file:

```
#####
```

```
# CACHE / HISTORIAN SETTINGS
```

```
sl.rtvadm.sub=$ftlServerRowExpirationTime:10
```

```
sl.rtvadm.sub=$ftlRowExpirationTimeForDelete:0
```

**Timestamp** The date and time this row of data was last updated in RTView.

### All Satellites Table

View the status, ID, and parent realm of all satellites. Each row in the table is a different satellite, and clicking on a table row drills-down to the "[Server Summary](#)" display allowing you to view performance metrics for the server hosting the satellite. You can click a column header to sort column data in ascending or descending order.

Server: Realm1

Count: 2

All Satellites Table

Parent Realm	Satellite	Satellite Label	Satellite Status	Satellite UUID
Realm1	Realm1_Sat1	192.168.200.229:8050	Running	d28a50aa-a0a8-44f0-b3ea-8da9359
Realm1	Realm1_Sat2	192.168.200.229:8050	Running	d28a50aa-a0a8-44f0-b3ea-8da9359

**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu Table open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

---

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by TIBCO FTL. Refer to the TIBCO FTL documentation for more information regarding these fields.

---

**Filter By:**

**Server** Select the server for which you want to view data.

**Count** The total number of table of satellite servers associated with the selected server. If the satellite server is not running, the row background is shown in yellow. When the satellite server is inactive (**Expired=true**) the row is shown in dark gray.

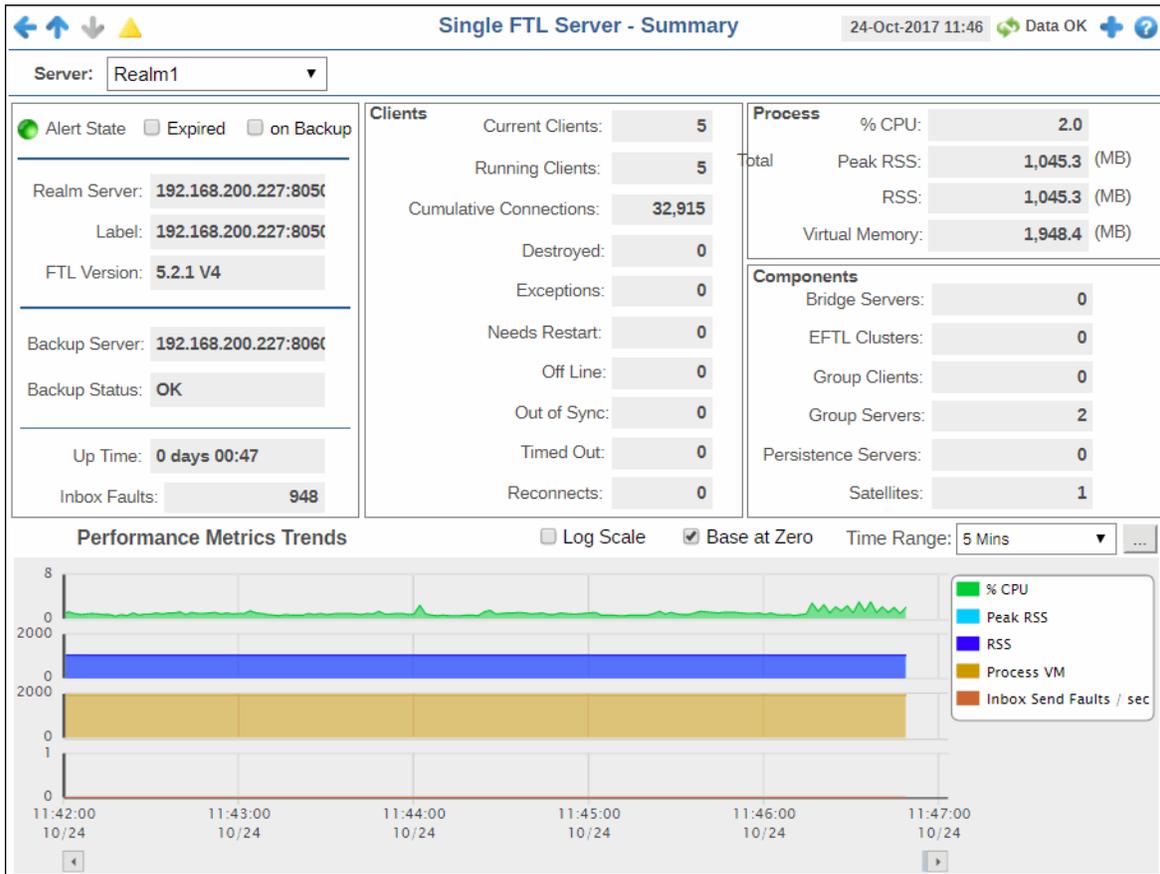
**All Satellites Table**

**Parent Realm** The name of the host server for the satellite.\*

<b>Satellite</b>	The satellite IP address and port number.*
<b>Satellite Status</b>	The server state. Valid values are:* <b>Running</b> -- the satellite is up and running. <b>Stopped</b> -- the satellite is stopped.
<b>Satellite UUID</b>	The unique ID for the satellite.*
<b>Expired</b>	<p>When checked (Expired=true), monitoring data for the row has not been received within the time specified for expiration, which is defined (in seconds) using the <b>\$ftlServerRowExpirationTime</b> substitution located in the <b>conf\rtvapm_tftlmon.properties</b> file. If the row has been expired for an extended period of time, the <b>\$ftlRowExpirationTimeForDelete</b> substitution determines when the row will be deleted from the cache that drives the display. The default values for the substitutions are 10 and 3600 seconds respectively, meaning that each of the rows will have Expired set to true after 10 seconds of inactivity and that expired FTL Server rows will be removed from the cache after 3600 seconds (one hour) of inactivity.</p> <p>To edit the default/current values, copy the lines below from <b>rtvapm_tftlmon.properties</b> file, paste them into the <b>sample.properties</b> file, and modify the lines in the <b>sample.properties</b> file:</p> <pre>##### # CACHE / HISTORIAN SETTINGS  sl.rtvapm.sub=\$ftlServerRowExpirationTime:10 sl.rtvapm.sub=\$ftlRowExpirationTimeForDelete:0</pre>
<b>Timestamp</b>	The date and time this row of data was last updated in RTView.

## Server Summary

This display allows you to investigate performance issues for the selected server. You can track current and historical performance metrics for a single FTL server and view how many components (satellites, EFTL clusters; bridge, group and persistence servers) the server hosts.



**Title Bar (possible features are):**

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu Table open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by TIBCO FTL. Refer to the TIBCO FTL documentation for more information regarding these fields.

**Filter By:**

**Server** Select the FTL Server for which you want to view data.

**Server Metrics**

**Alert State** The maximum alert level on the server:

- One or more alerts have exceeded their specified **ALARMLEVEL** threshold.
- One or more alerts have exceeded their specified **WARNINGLEVEL** threshold.
- No alerts have exceeded an alert threshold.

**Expired** When checked (Expired=true), monitoring data for the FTL Server row has not been received within the time specified for expiration, which is defined (in seconds) using the **\$ftlServerRowExpirationTime** substitution located in the **conf\rtvapm\_tftlmon.properties** file. If the row has been expired for an extended period of time, the **\$ftlRowExpirationTimeForDelete** substitution determines when the row will be deleted from the cache that drives the display. The default values for the substitutions are 10 and 3600 seconds respectively, meaning that each of the FTL Server rows will have Expired set to true after 10 seconds of inactivity and that expired FTL Server rows will be removed from the cache after 3600 seconds (one hour) of inactivity.

To edit the default/current values, copy the lines below from **rtvapm\_tftlmon.properties** file, paste them into the **sample.properties** file, and modify the lines in the **sample.properties** file:

```
#####
# CACHE / HISTORIAN SETTINGS

sl.rtvapm.sub=$ftlServerRowExpirationTime:10
sl.rtvapm.sub=$ftlRowExpirationTimeForDelete:0
```

**on Backup** When checked, this server has a backup server.\*

**Realm Server** The server IP address or host name.\*

**Label** The server label.\*

**FTL Version** The FTL version on the server.\*

**Backup Server** The IP address and port of the backup server.\*

**Backup Status** The current backup server status.\*

**Up Time** The number of days, hours and minutes since the server was last started.\*

**Inbox Faults** The total number of faults when sending messages to inbox subscribers.\*

**Satellites** The number of satellites.\*

**Clients**

**Current Clients** The number of clients currently on the server.\*

**Running Clients** The number of clients currently active on the server.\*

<b>Cumulative Client Connections</b>	The total number of clients the server has connected since it the server was last started.*
<b>Destroyed</b>	The total number of destroyed clients since the server was last started.*
<b>Exceptions</b>	The total number of client exceptions since the server was last started.*
<b>Needs Restart</b>	The total number of clients that had to reconnect since the server was last started.*
<b>Offline</b>	The number of clients currently offline.*
<b>Out of Sync</b>	The number of clients currently out of sync.*
<b>Timed out</b>	The total number of clients connections that timed out due to inactivity since the server was last started.*
<b>Reconnects</b>	The total number of clients that had to reconnect since the server was last started.*

**Process**

<b>% CPU</b>	The amount of CPU used, in percent.*
<b>Peak RSS</b>	The maximum RSS memory used, in kilobytes.*
<b>RSS</b>	The current RSS memory being used, in megabytes.*
<b>Virtual Memory</b>	The current virtual memory being used, in megabytes.*

**Components**

<b>Bridge Servers</b>	The number of bridge servers connected.*
<b>EFTL Clusters</b>	The number of EFTL clusters connected.*
<b>Group Clients</b>	The number of group clients connected.*
<b>Group Servers</b>	The number of group servers connected.*
<b>Persistence Servers</b>	The number of persistence servers connected.*
<b>Satellites</b>	The number of satellites connected.*

**Performance Metrics Trends**

Traces the following for the selected server:

**% CPU** -- The percent CPU used.

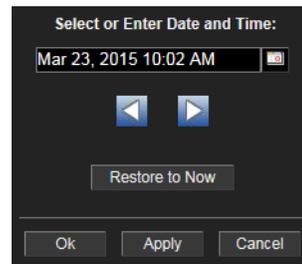
**Peak RSS** -- The maximum amount of RSS memory used.

**RSS** -- The RSS memory used.

**Process VM** -- The current virtual memory used.

**Inbox Send Faults/sec** -- The rate of faults when sending messages to inbox subscribers (per second).

- Log Scale** This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.
- Base at Zero** When this option is checked, zero is set as the Y axis minimum for all graph traces.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## FTL Clients

These displays present performance metrics and alert status for FTL clients. Available displays are:

- **"All Clients Heatmap"**: Shows alert status for all FTL clients in all realms in a heatmap format.
- **"All Clients Table"**: Shows all available utilization metrics for all FTL clients on a selected server in a tabular format.
- **"Clients by Group"**: Lists all FTL clients by group with member details and CPU utilization in a tabular format.
- **"Client Summary"**: Displays current and historical metrics for a single FTL client.
- **"Client Metrics"**: Provides details about metric calculations for each FTL client.

## All Clients Heatmap

This heatmap allows you to view the status and alerts of all FTL clients. You can view the clients in the heatmap based on the following metrics: the current alert severity, the current alert count, the current CPU usage, the rate of messages received, and the rate of messages sent.

Each heatmap rectangle represents a client, and clicking on a client drills-down to the “[Client Summary](#)” display allowing you to view metrics for that client. The rectangle color indicates the most critical alert state. Clicking on the  icon in the upper left hand corner of the display toggles between the commonly accessed **Table** and **Heatmap** displays. You can also mouse-over the rectangles to view more details about host performance and status.



### Title Bar (possible features are):

-   Open the previous and upper display.
-  Open an instance of this display in a new window.
-  Open the online help page for this display.
-   open commonly accessed displays.

 **Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

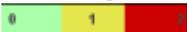
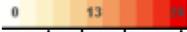
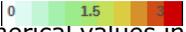
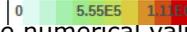
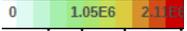
 **23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

### Filter By:

**Server** Choose the server for which you want to view data.

### Fields and Data

This display includes:

<b>Names</b>	Select this check box to display the names of clients.
<b>Log</b>	This option enables visualization on a logarithmic scale, and should be used when the range in your data is very broad. For example, if you have data that ranges from the tens to the thousands, then data in the range of tens will be neglected visually if you do not check this option. This option makes data on both extreme ranges visible by using the logarithmic of the values rather than the actual values.
<b>Auto</b>	When checked, the values of the selected metric are auto-scaled to its highest defined value. When unchecked, the values of the selected metric display based on the threshold defined for the alert associated with the selected metric. Selecting Auto helps to visualize the range of the values currently present for the selected metric instead of the threshold of the alert that has been associated with the metric. All metrics that have not been associated in the heatmap defaults with alerts use a monochromatic color gradient bar (whites and greens). All metrics that have been associated in the heatmap defaults with alerts use a multi-chromatic color gradient bar (reds, yellows, white, and greens).
<b>Metric</b>	Select the metric driving the heatmap display. The default is Alert Severity. Each <b>Metric</b> has a color gradient bar that maps values to colors. Each rectangle represents a client. Mouse-over any rectangle to display the current values of the metrics for the client. Click on a rectangle to drill-down to the associated " <a href="#">Client Summary</a> " display for a detailed view of metrics for that particular client.
<b>Alert Severity</b>	<p>The maximum alert level in the item (index) associated with the rectangle. Values range from <b>0</b> to <b>2</b>, as indicated in the color gradient bar , where <b>2</b> is the greatest <b>Alert Severity</b>.</p> <p><b>2</b> Metrics that have exceeded their specified <b>ALARMLEVEL</b> threshold and have an Alert Severity value of <b>2</b> are shown in red. For a given rectangle, this indicates that one or more metrics have exceeded their alarm threshold.</p> <p><b>1</b> Metrics that have exceeded their specified <b>WARNINGLEVEL</b> threshold and have an Alert Severity value of <b>1</b> are shown in yellow. For a given rectangle, this indicates that one or more metrics have exceeded their warning threshold.</p> <p><b>0</b> Metrics that have not exceeded either specified threshold have an Alert Severity value of <b>0</b> and are shown in green. For a given rectangle, this indicates that no metrics have exceeded a specified alert threshold.</p>
<b>Alert Count</b>	<p>The total number of alarm and warning alerts in a given item (index) associated with the rectangle.</p> <p>The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>CPU Usage</b>	<p>The total amount of CPU used. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the defined alert threshold of <b>TftlClientCpuUsageHigh</b>. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>Memory</b>	<p>The current amount of memory used for processing. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the defined alert threshold of <b>TftlClientMemoryHigh</b>. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>V(irtual) Memory</b>	<p>The current amount of virtual memory being used for processing. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the defined alert threshold of <b>TftlClientVirtualMemoryHigh</b>. The middle value in the gradient bar indicates the middle value of the range.</p>

**Msgs Rcvd/sec** The rate of messages received (per second). The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the defined alert threshold of **TftlClientMsgsRcvdRateHigh**. The middle value in the gradient bar indicates the middle value of the range.

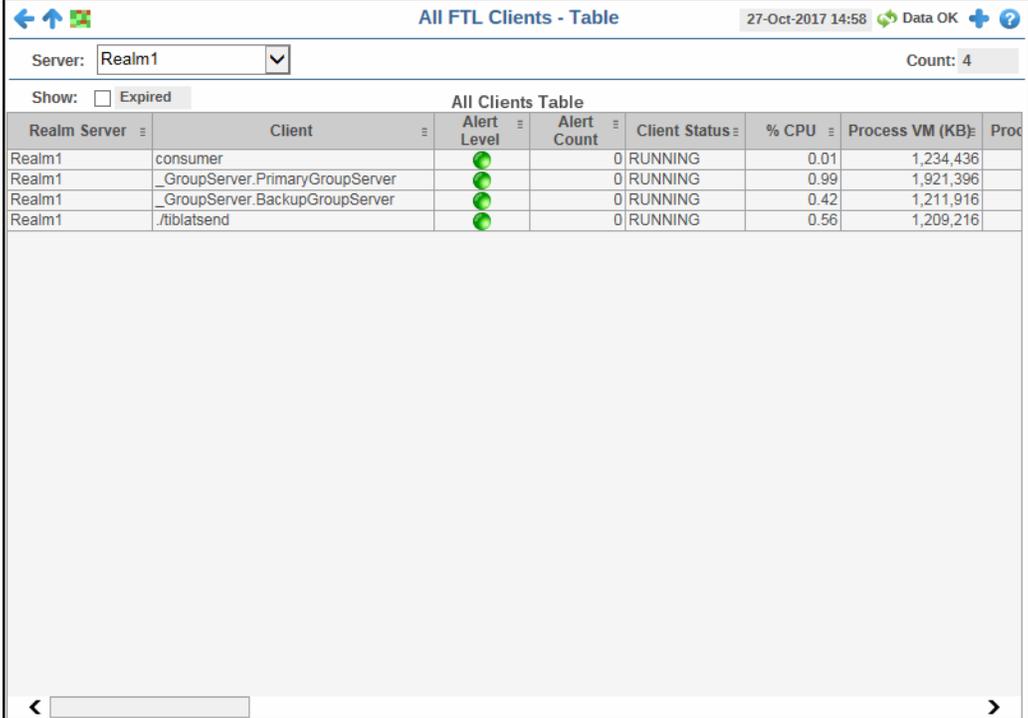
**Msgs Sent/sec** The rate of messages sent (per second). The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the defined alert threshold of **TftlClientMsgsSentRateHigh**. The middle value in the gradient bar indicates the middle value of the range.

## All Clients Table

Investigate detailed utilization metrics for all FTL clients. The **All Clients Table** contains all metrics available for clients, including the number of current client connections. Each row in the table contains data for a particular client, and clicking on a table row drills-down to the “[Client Summary](#)” display allowing you to view metrics for that particular client. You can click a column header to sort column data in ascending or descending order. Clicking the  icon in the upper left hand corner of the display toggles between the commonly accessed **Table** and **Heatmap** displays.

Note that the number of clients found in this table might not match the client count found in the “[All Servers Table](#)”, possibly due to the following:

- One client might have one or more group joins resulting in a higher client count. For example, if a client has two group joins, CLIENT\_COUNT equals 3, but will only be listed as a single client in this table.
- A TIBCO bridge could have one or more logical bridges running inside the bridge process, which could result in an increased CLIENT\_COUNT even though there is actually only one client.
- Other FTL services could get a Client ID and, hence, be included in the CLIENT\_COUNT even though they are not necessarily clients.



Realm Server	Client	Alert Level	Alert Count	Client Status	% CPU	Process VM (KB)	Proc
Realm1	consumer		0	RUNNING	0.01	1,234,436	
Realm1	_GroupServer.PrimaryGroupServer		0	RUNNING	0.99	1,921,396	
Realm1	_GroupServer.BackupGroupServer		0	RUNNING	0.42	1,211,916	
Realm1	.tiblatsend		0	RUNNING	0.56	1,209,216	



**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by TIBCO FTL. Refer to the TIBCO FTL documentation for more information regarding these fields.

### Filter By:

**Server** Select the server for which you want to view data.

**Show: Expired** Select this check box to display those rows containing expired data. Leave unchecked to display only those rows that are not expired.

**Count** The total number of active, inactive, and standby FTL clients. **Inactive Clients** are shown in dark red. **Standby Clients** are shown in blue.

### All Clients Table

**Realm Server** The name of the server.

**Client** The name of the client.

**Alert Level** The maximum alert level. Values range from 0 to 2, as indicated in the color gradient bar, where **2** is the greatest Alert Severity.

 One or more alerts have exceeded their specified **ALARMLEVEL** threshold, have an Alert Severity value of **2**, and are shown in red.

 One or more alerts have exceeded their specified **WARNINGLEVEL** threshold, have an Alert Severity value of **1**, and are shown in yellow.

 No alerts have exceeded an alert threshold, which have an Alert Severity value of **0**, and are shown in green.

**Alert Count** The current number of alerts.

**Client Status** The status of the client. For example, RUNNING.\*

**% CPU** The amount of CPU used, in percent.\*

**Process VM (KB)** The amount of virtual memory used for processing, in kilobytes. \*

**Process RSS (KB)** The current RSS memory being used, in kilobytes.\*

**Process Peak RSS (KB)** The maximum RSS memory used, in kilobytes.\*

**Msgs Rcvd/ sec** The number of messages received per second.

<b>Msgs Sent/sec</b>	The number of messages sent per second.
<b>Delta Msgs Rcvd</b>	The total number of messages received since the last data update.
<b>Delta Msgs Sent</b>	The total number of messages sent since the last data update.
<b>Total Msgs Rcvd</b>	The total number of messages received since the client started.*
<b>Total Msgs Sent</b>	The total number of messages sent since the client started.*
<b>Store Mismatch Msgs</b>	Any non-zero value indicates a store mismatch misconfiguration, which occurs when a direct path transport connects two endpoints that are associated with two different persistence stores.*
<b>Dynamic Formats</b>	The number of distinct dynamic formats that the client creates within the sample interval.*
<b>User CPU</b>	The amount of time the user has used the CPU, in microseconds.*
<b>Application</b>	The application name of the client.*
<b>Application Instance</b>	The application instance of the client.*
<b>Client ID</b>	The unique ID for the client.*
<b>Process ID</b>	The unique ID for the process.*
<b>FTL User</b>	The FTL user that is being used by the client.*
<b>Effective User</b>	The UID of the client (which is used for most access checks).*
<b>Host</b>	The host name.*
<b>Host IP</b>	The host IP address.*
<b>FTL Version</b>	The FTL version on the host.*
<b>Expired</b>	<p>When checked (Expired=true), monitoring data for the row has not been received within the time specified for expiration, which is defined (in seconds) using the <b>\$fttlRowExpirationTime</b> substitution located in the <b>conf\rtvapm_tftlmon.properties</b> file. If the row has been expired for an extended period of time, the <b>\$fttlRowExpirationTimeForDelete</b> substitution determines when the row will be deleted from the cache that drives the display. The default values for the substitutions are 120 and 3600 seconds respectively, meaning that each of the rows will have Expired set to true after 120 seconds of inactivity and that expired rows will be removed from the cache after 3600 seconds (one hour) of inactivity.</p> <p>To edit the default/current values, copy the lines below from <b>rtvapm_tftlmon.properties</b> file, paste them into the <b>sample.properties</b> file, and modify the lines in the <b>sample.properties</b> file:</p> <pre>##### # CACHE / HISTORIAN SETTINGS  sl.rtvview.sub=\$fttlServerRowExpirationTime:120 sl.rtvview.sub=\$fttlRowExpirationTimeForDelete:0</pre>

**Data Timestamp** The date and time this row of data occurred in FTL.\*

**Local Timestamp** The date and time this row of data was last updated in RTView.

## Clients by Group

This display lists all clients and their associated FTL groups for a particular server. Each row in the table is a different client, and clicking on a table row drills-down to the "Client Summary" display allowing you to view metrics for that particular client. You can click a column header to sort column data in ascending or descending order.

Group Name	Client Label	Ordinal	Member Type	Client Status	% CPU	Client ID
testgroup1		1	FULL_MEMBER			
testgroup1		0	OBSERVER_MEMBER			
testgroup1		2	FULL_MEMBER			
testgroup3		0	OBSERVER_MEMBER			
testgroup3		1	FULL_MEMBER			
testgroup3		2	FULL_MEMBER			
testgroup4		1	FULL_MEMBER			
testgroup4		2	FULL_MEMBER			
testgroup4		0	OBSERVER_MEMBER			

### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by TIBCO FTL. Refer to the TIBCO FTL documentation for more information regarding these fields.

**Filter By:**

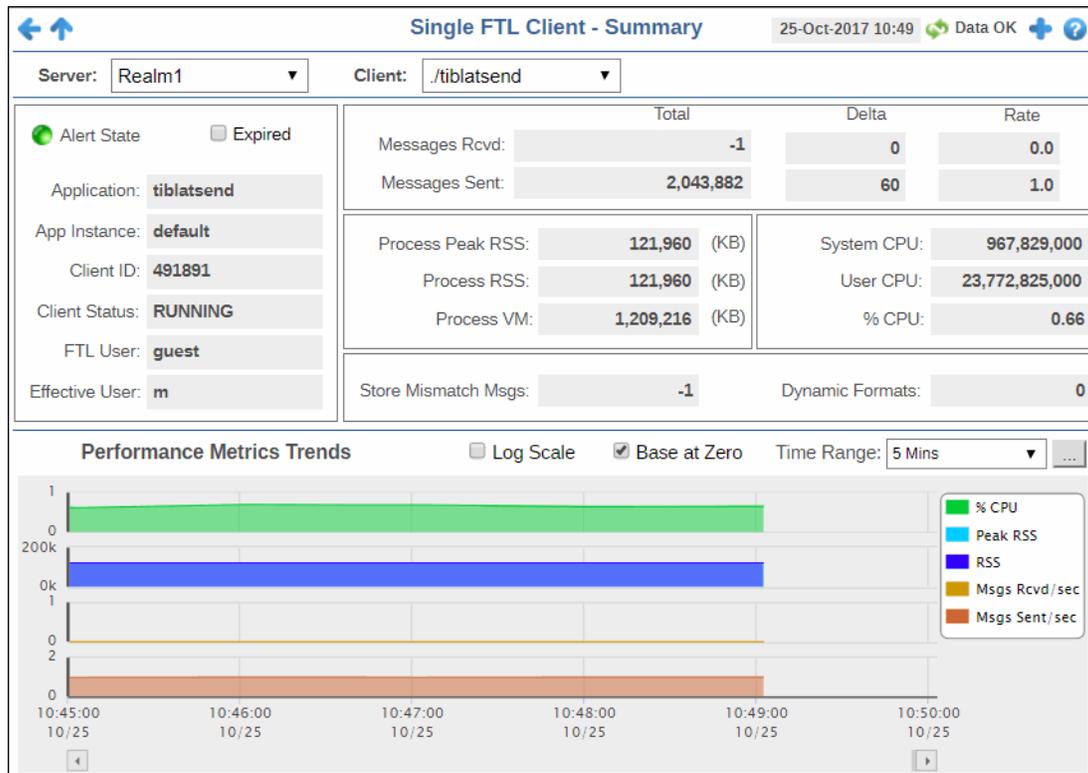
<b>Server</b>	Select the server for which you would like to view data.
<b>Group Server Mode</b>	The server mode. Valid values are: <b>Primary</b> -- The server is acting as a primary server. <b>Secondary</b> --The the server is acting as a secondary server.
<b>State</b>	<b>Running</b> -- The client is running and its local realm definition is up-to-date. <b>Needs Restart</b> – The client needs to be restarted to update its local realm definition. <b>Timed Out</b> - The server has lost the heartbeat signal from the client. Either the client has stopped or a network issue is obstructing the signal. <b>Exception</b> - The client is running, but its realm definition is out-of-date. <b>Out-of-Sync</b> -- The realm definition revision on the client and server are different.
<b>Count</b>	The total number of FTL clients. Inactive clients are shown in dark gray and when the client is associated with an active backup server the client is are shown in yellow.

**Clients by Group Table**

<b>Group Name</b>	The name of the group.*
<b>Client Label</b>	The client's label.*
<b>Ordinal</b>	The number representing the client's position within the group. A value of -1 indicates that the client has been disconnected from the group server.*
<b>Member Type</b>	The client's member type.*
<b>Client Status</b>	The status of the client. For example, RUNNING.*
<b>% CPU</b>	The amount of CPU used, in percent.*
<b>Client ID</b>	The ID of the client.*
<b>Process ID</b>	The process ID.*
<b>Expired</b>	When checked (Expired=true), monitoring data for the row has not been received within the time specified for expiration, which is defined (in seconds) using the <b>\$ftlRowExpirationTime</b> substitution located in the <b>conf\rtvadm_tftlmon.properties</b> file. If the row has been expired for an extended period of time, the <b>\$ftlRowExpirationTimeForDelete</b> substitution determines when the row will be deleted from the cache that drives the display. The default values for the substitutions are 120 and 3600 seconds respectively, meaning that each of the rows will have Expired set to true after 120 seconds of inactivity and that expired rows will be removed from the cache after 3600 seconds (one hour) of inactivity.  To edit the default/current values, copy the lines below from <b>rtvadm_tftlmon.properties</b> file, paste them into the <b>sample.properties</b> file, and modify the lines in the <b>sample.properties</b> file: <pre>##### # CACHE / HISTORIAN SETTINGS  sl.rtvadm.sub=\$ftlServerRowExpirationTime:120 sl.rtvadm.sub=\$ftlRowExpirationTimeForDelete:0</pre>
<b>Timestamp</b>	The date and time this row of data was last updated in RTView.

## Client Summary

Track current and historical performance metrics for a single FTL client. Use this display to investigate performance issues of a client.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by TIBCO FTL. Refer to the TIBCO FTL documentation for more information regarding these fields.

### Filter By:

- Server** Choose an FTL server to display.
- Client** Choose a client to display.

### Fields and Data

**Alert State** The maximum alert level has been exceeded. Values range from 0 to 2, as indicated in the color gradient bar, where **2** is the greatest Alert Severity.

- One or more alerts have exceeded their specified **ALARMLEVEL** threshold, have an Alert Severity value of **2**, and are shown in red.
- One or more alerts have exceeded their specified **WARNINGLEVEL** threshold, have an Alert Severity value of **1**, and are shown in yellow.
- No alerts have exceeded an alert threshold, which have an Alert Severity value of **0**, and are shown in green.

**Expired** When checked (Expired=true), monitoring data for the row has not been received within the time specified for expiration, which is defined (in seconds) using the **\$ftlRowExpirationTime** substitution located in the **conf\rtvadm\_tftlmon.properties** file. If the row has been expired for an extended period of time, the **\$ftlRowExpirationTimeForDelete** substitution determines when the row will be deleted from the cache that drives the display. The default values for the substitutions are 120 and 3600 seconds respectively, meaning that each of the rows will have Expired set to true after 120 seconds of inactivity and that expired rows will be removed from the cache after 3600 seconds (one hour) of inactivity.

To edit the default/current values, copy the lines below from **rtvadm\_tftlmon.properties** file, paste them into the **sample.properties** file, and modify the lines in the **sample.properties** file:

```
#####
# CACHE / HISTORIAN SETTINGS

sl.rtvadm.sub=$ftlServerRowExpirationTime:120
sl.rtvadm.sub=$ftlRowExpirationTimeForDelete:0
```

**Application** The application name of the client.\*

**App Instance** The application instance of the client.\*

**Client ID** The unique identifier for the client.\*

**Client Status** The status of the client. For example, RUNNING.\*

**FTL User** The FTL user that is being used by the client.\*

**Effective User** The UID of the client (which is used for most access checks).\*

**Messages Rcvd**

- Total** The total number of messages received since the client started.\*
- Delta** The number of messages received since the last data update.
- Rate** The number of messages received per second.

**Messages Sent**

- Total** The total number of messages sent since the client started.\*
- Delta** The number of messages sent since the last data update.
- Rate** The number of messages sent per second.

**Process Peak RSS** The maximum RSS memory used, in kilobytes.\*

**Process RSS** The current RSS memory being used, in kilobytes.\*

**Process VM** The current virtual memory being used, in kilobytes.\*

<b>System CPU</b>	The amount of CPU used by the system, in kilobytes.*
<b>User CPU</b>	The amount of CPU used by the client, in kilobytes.*
<b>% CPU</b>	The percent of CPU used by the client.*
<b>Store Mismatch Msgs</b>	Any non-zero value indicates a store mismatch misconfiguration, which occurs when a direct path transport connects two endpoints that are associated with two different persistence stores.*
<b>Dynamic Formats</b>	The number of distinct dynamic formats that the client creates within the sample interval.*

### Performance Metrics Trends

Traces the following for the selected client:

**% CPU** -- The amount of CPU used, in percent.

**Peak RSS** -- The maximum RSS memory used, in kilobytes.

**RSS** -- The current RSS memory being used, in kilobytes.

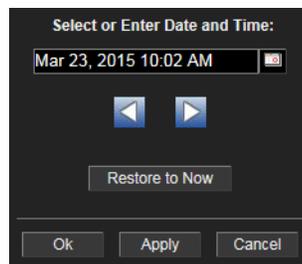
**Msgs Rcvd/sec** -- The number of messages received per second.

**Msgs Sent/sec** -- The number of messages sent per second.

**Log Scale** This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.

**Base at Zero** When this option is checked, zero is set as the Y axis minimum for all graph traces.

**Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



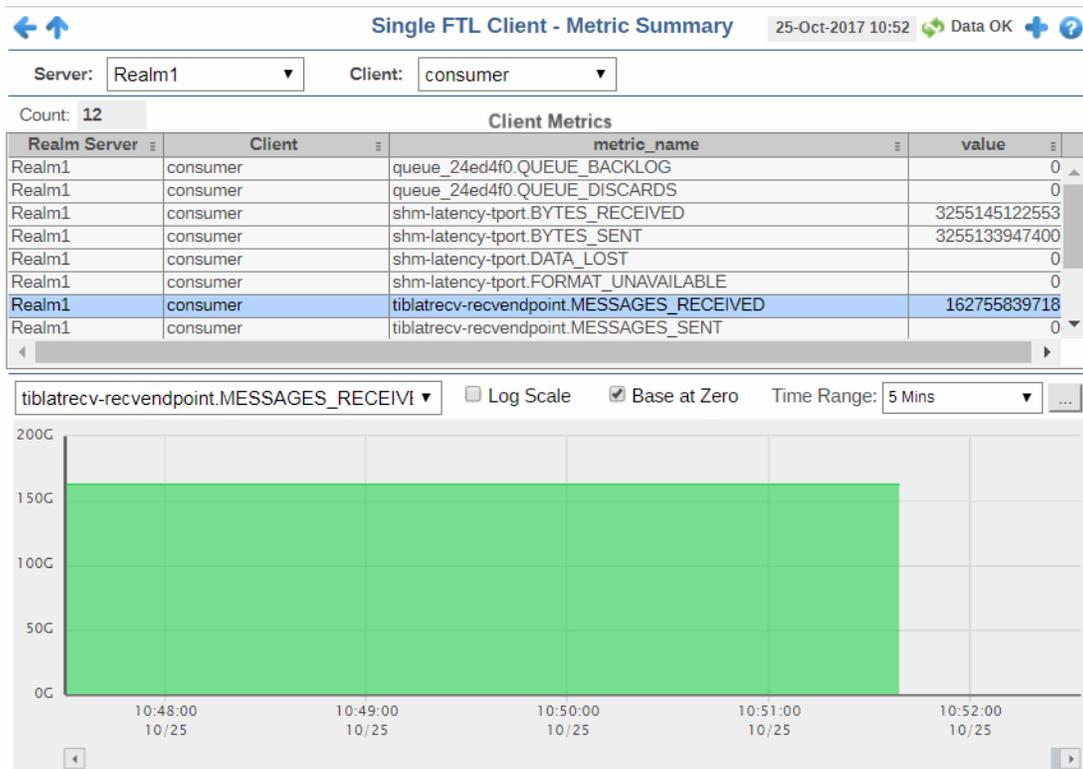
By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Client Metrics

Track detailed performance and utilization metrics for a single FTL client.



**Title Bar** (possible features are):

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu Table open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by TIBCO FTL. Refer to the TIBCO FTL documentation for more information regarding these fields.

**Filter By:**

- Server** Select the server containing the client for which you want to see data.
- Client** Select the client for which you want to see data.
- Count** The number of rows/metrics in the table.

### Client Metrics Table

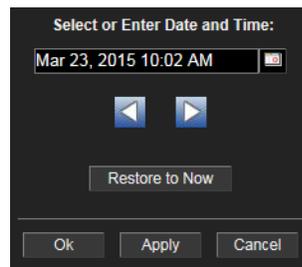
<b>Realm Server</b>	The name of the server.
<b>Client</b>	The name of the client.
<b>metric_name</b>	The name of the metric.*
<b>value</b>	The current value of the metric.*

**Trend Graph** Select a metric from the drop-down menu to trace in the trend graph for the selected client.

**Log Scale** This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.

**Base at Zero** When this option is checked, zero is set as the Y axis minimum for all graph traces.

**Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## FTL Events

This display presents FTL events for a particular client. The available display is:

- **“Events”**: Table that lists all FTL events for a particular client.

### Events

This display allows you to view FTL events on one or all clients. You can view the event type, the client status, the associated application, and the FTL user name among other details.

**All FTL Events - Table** 25-Oct-2017 10:54 Data OK

Server:  Client:

Count:  Time Range:

Timestamp	Server	Client	Event Type	Client Status	Client
03-Oct-2017 06:32:16.801	Realm1	client41	DISCONNECTED	DESTROYED	1361
03-Oct-2017 06:32:16.801	Realm1	client41	DISCONNECTED	DESTROYED	1355
03-Oct-2017 06:32:16.801	Realm1	client42	DISCONNECTED	DESTROYED	1351
03-Oct-2017 06:32:16.801	Realm1	client43	DISCONNECTED	DESTROYED	1353
03-Oct-2017 06:32:16.801	Realm1	client43	DISCONNECTED	DESTROYED	1357

**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Note:** Fields/columns with an asterisk (\*) at the end of the field/column definition contain data that is provided by TIBCO FTL. Refer to the TIBCO FTL documentation for more information regarding these fields.

**Filter By:**

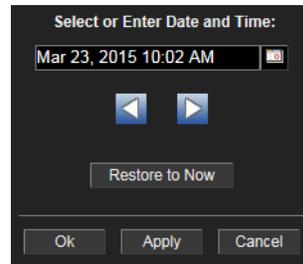
**Server** Select an FTL server containing the client for which you want to view data.

**Client** Select an FTL client for which you want to view data.

**Count** The number of rows in the table.

**Time Range**

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the  button.



By default, the time range end point is the current time. To change the time range end point, click the  button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. **Note:** The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

**Fields and Data**

<b>Timestamp</b>	The date and time this row of data was last updated in RTView.
<b>Server</b>	The name of the server.
<b>Client</b>	The name of the client.
<b>Event Type</b>	The type of event.*
<b>Client Status</b>	The status of the client.*
<b>Client ID</b>	The ID of the client.*
<b>Process ID</b>	The process ID.*
<b>Application</b>	The application name of the client.*
<b>Application Instance</b>	The name of the application instance.*
<b>FTL User</b>	The FTL user that is being used by the client.*
<b>Realm Server</b>	The IP address and port of the realm server.*
<b>Host</b>	The name of the host.*
<b>IP Address</b>	The IP address of the client.*
<b>Version</b>	The version of the client running on the server.*

---

## TIBCO FTL - HTML

The Solution Package for TIBCO FTL HTML displays provide extensive visibility into the health and performance of the TIBCO FTL servers. The HTML version features an overview display, "[TIBCO FTL Overview - HTML](#)" (shown below), and the following Views which can be found under **Components** tab > **Middleware**.

- "[FTL Servers - HTML](#)"
- "[FTL Clients - HTML](#)"

### TIBCO FTL Overview - HTML

The **TIBCO FTL Overview** is the top-level display for the TIBCO FTL Solution Package, which provides a good starting point for immediately getting the status of all your TIBCO FTL servers on your Data Server.

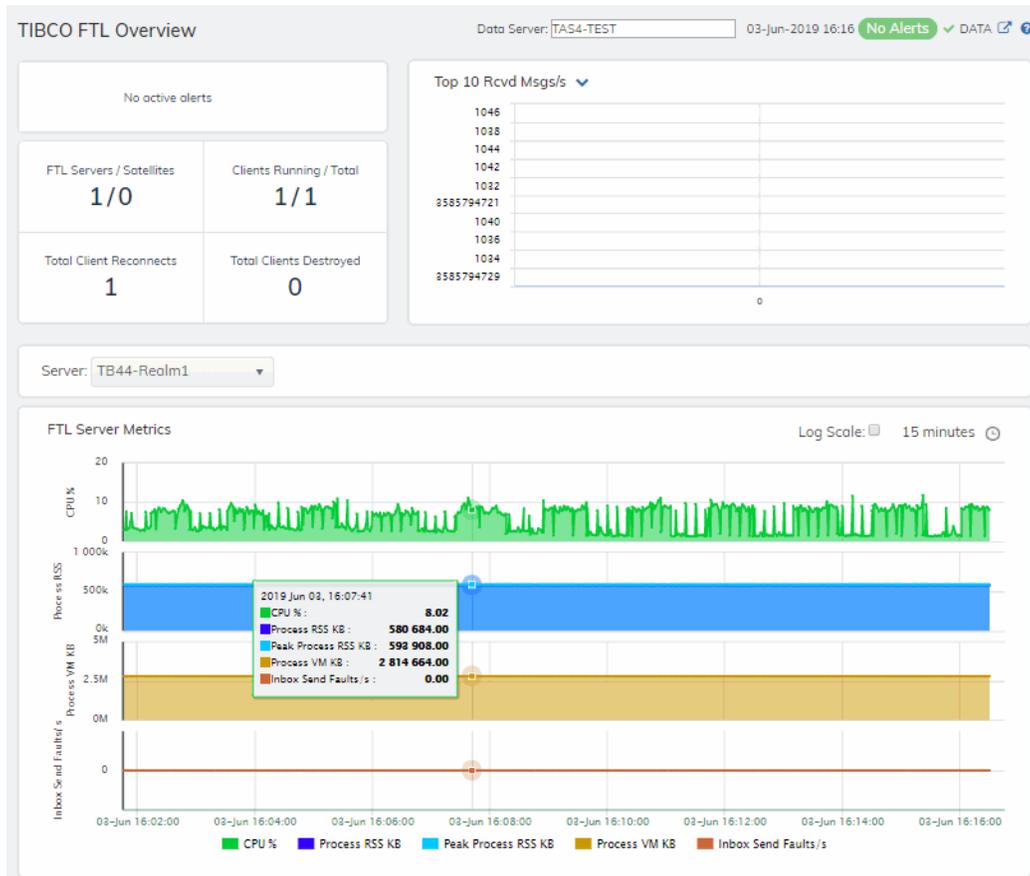
You can select the RTView Data Server for which you want to see data and easily view the current data for that Data Server including:

- The total number of active alerts, including the total number of critical and warning alerts.
- The number of **FTL Servers** and **Satellites** across all servers.
- The number of **Client Reconnects** and **Clients Destroyed** across all servers.
- A bar graph of the servers with **Top 10 Rcvd Msgs/s**. You can also select **Top 10 Sent Msgs/s**, **Top 10 Timed Out Clients**, **Top 10 Client Rcvd Msgs/s**, **Top 10 Client Sent Msgs/s** or **Top 10 Server CPU%**.

You can hover over each region in the upper half of the Overview to see more detail. You can also drill down to see even more detail by clicking on each respective region in the Overview. For example, clicking on **FTL Servers** / **Satellites** opens the "[FTL Servers Table - HTML](#)" display.

The bottom half of the display provides a performance trend graph that traces **CPU%**, **Process RSS KB**, **Peak Process RSS KB**, **Process VM KB** and **Inbox Send Faults/s** for a selected server.

You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.



## FTL Servers - HTML

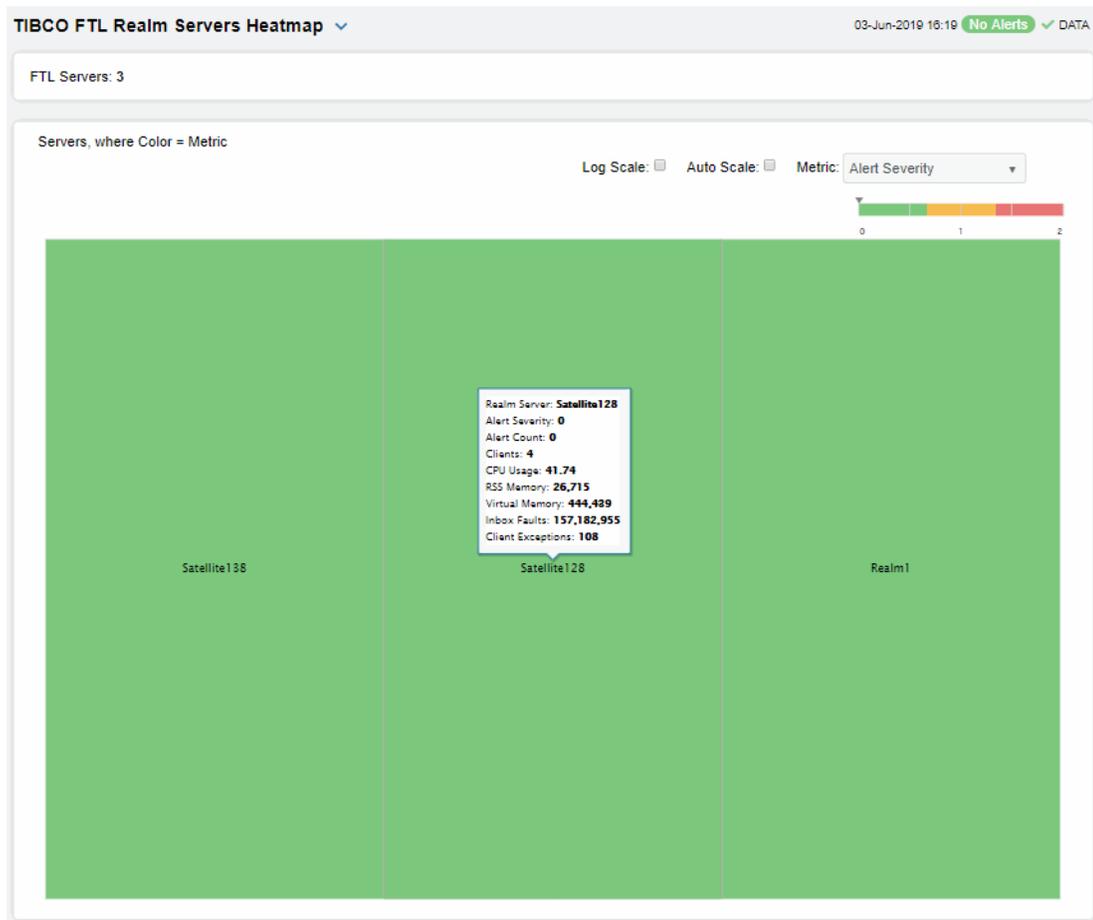
These displays present performance metrics and alert status for all FTL servers. Displays are:

- [“FTL Servers Heatmap - HTML”](#): Heatmap shows server and alert status for all FTL servers in all realms.
- [“FTL Servers Table - HTML”](#): Table shows all available utilization metrics for all FTL servers.
- [“FTL Server Summary - HTML”](#): Current and historical metrics for a single FTL server.

## FTL Servers Heatmap - HTML

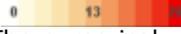
This heatmap display allows you to quickly identify the current status of each of your servers for each available metric. You can view the servers in the heatmap based on the following metrics: the current **Alert Severity**, the current **Alert Count**, **Clients** (the total number of clients), **CPU Usage** (the current amount of CPU being used), **RSS Memory** (the current amount of memory being used for processing), **Virtual Memory** (the current amount of virtual memory being used for processing) and **Inbox Faults** (the number of inbox send faults) and **Client Exceptions**. By default, this display shows the heatmap based on the **Alert Severity** metric.

Each heatmap rectangle represents a server. The rectangle color indicates the most critical alert state. You can click on a rectangle to drill-down to the "[FTL Server Summary - HTML](#)" display and view metrics for that server. Clicking on the  icon next to the display title toggles between the commonly accessed **Table** and **Heatmap** displays. Mouse-over the rectangles to view more details about host performance and status. You can view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.

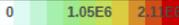


## Fields and Data

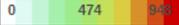
This display includes:

<b>Names</b>	Select to display the names of servers on the hosts.
<b>Metric</b>	Select the metric driving the heatmap display. The default is Alert Severity. Each <b>Metric</b> has a color gradient bar that maps values to colors. The heatmap organizes the servers by host, where each rectangle represents a server. Mouse-over any rectangle to display the current values of the metrics for the Server. Click on a rectangle to drill-down to the associated <a href="#">"FTL Server Summary - HTML"</a> display for a detailed view of metrics for that particular server.
<b>Alert Severity</b>	<p>The maximum alert level in the item (index) associated with the rectangle. Values range from <b>0</b> to <b>2</b>, as indicated in the color gradient bar , where <b>2</b> is the greatest <b>Alert Severity</b>.</p> <p><b>2</b> Metrics that have exceeded their specified <b>ALARMLEVEL</b> threshold and have an Alert Severity value of <b>2</b> are shown in red. For a given rectangle, this indicates that one or more metrics have exceeded their alarm threshold.</p> <p><b>1</b> Metrics that have exceeded their specified <b>WARNINGLEVEL</b> threshold and have an Alert Severity value of <b>1</b> are shown in yellow. For a given rectangle, this indicates that one or more metrics have exceeded their warning threshold.</p> <p><b>0</b> Metrics that have not exceeded either specified threshold have an Alert Severity value of <b>0</b> and are shown in green. For a given rectangle, this indicates that no metrics have exceeded a specified alert threshold.</p>
<b>Alert Count</b>	<p>The total number of alarm and warning alerts in a given item (index) associated with the rectangle.</p> <p>The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>Client Count</b>	<p>The total number of clients in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the defined alert threshold of <b>TftlServerClientCountHigh</b>. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>CPU Usage</b>	<p>The total amount of CPU used. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the defined alert threshold of <b>TftlServerCpuUsageHigh</b>. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>Memory</b>	<p>The current memory being used. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the defined alert threshold of <b>TftlServerMemoryHigh</b>. The middle value in the gradient bar indicates the middle value of the range.</p>

**Virtual Memory**

The current virtual memory being used. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the defined alert threshold of **TftlServerVirtualMemoryHigh**. The middle value in the gradient bar indicates the middle value of the range.

**#Inbox Faults**

The total number of inbox faults. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the defined alert threshold of **TftlServerInboxSendFaultsHigh**. The middle value in the gradient bar indicates the middle value of the range.

**FTL Servers Table - HTML**

Investigate detailed utilization metrics and configuration settings for all FTL servers. The **TIBCO FTL Servers Table** contains all metrics available for servers, including the number of current client connections.

Each row in the table contains data for a particular **Realm Server**. You can search, filter, sort and choose columns to include by clicking a column header icon (to the right of each column label) and selecting **Filter**, **Sort Ascending**, **Sort Descending** or **Columns**. Or just click a column header to sort.

Right-click on a table cell to **Export to Excel** or **Copy Cell Value**.

You can click on a row to drill-down to the ["FTL Server Summary - HTML"](#) display and view details for that server.

TIBCO FTL Servers Table 03-Jun-2019 16:21 No Alerts ✓ DATA

FTL Servers: 3

Realm Server	Realm Server Label	Alert Level	Alert Count	Backup Server Is Active	Backup Server Status	CPU %	Clients	Clients Running	Cumulative Client Connects
Satellite138	192.168.1.138:8050	✓	0		OK	58.15	6	5	
Satellite128	192.168.1.128:8050	✓	0		OK	41.85	6	4	
Realm1	192.168.1.149:8050	✓	0		OK	42.58	5	5	

## FTL Server Summary - HTML

Track current and historical utilization and performance metrics for a selected FTL server. You can view how many components (such as satellites, EFTL clusters, bridges, groups and persistence servers) the server hosts.

Clicking on the metric boxes at the top of the display takes you to the "[FTL Servers Table - HTML](#)" display, where you can view additional server data.

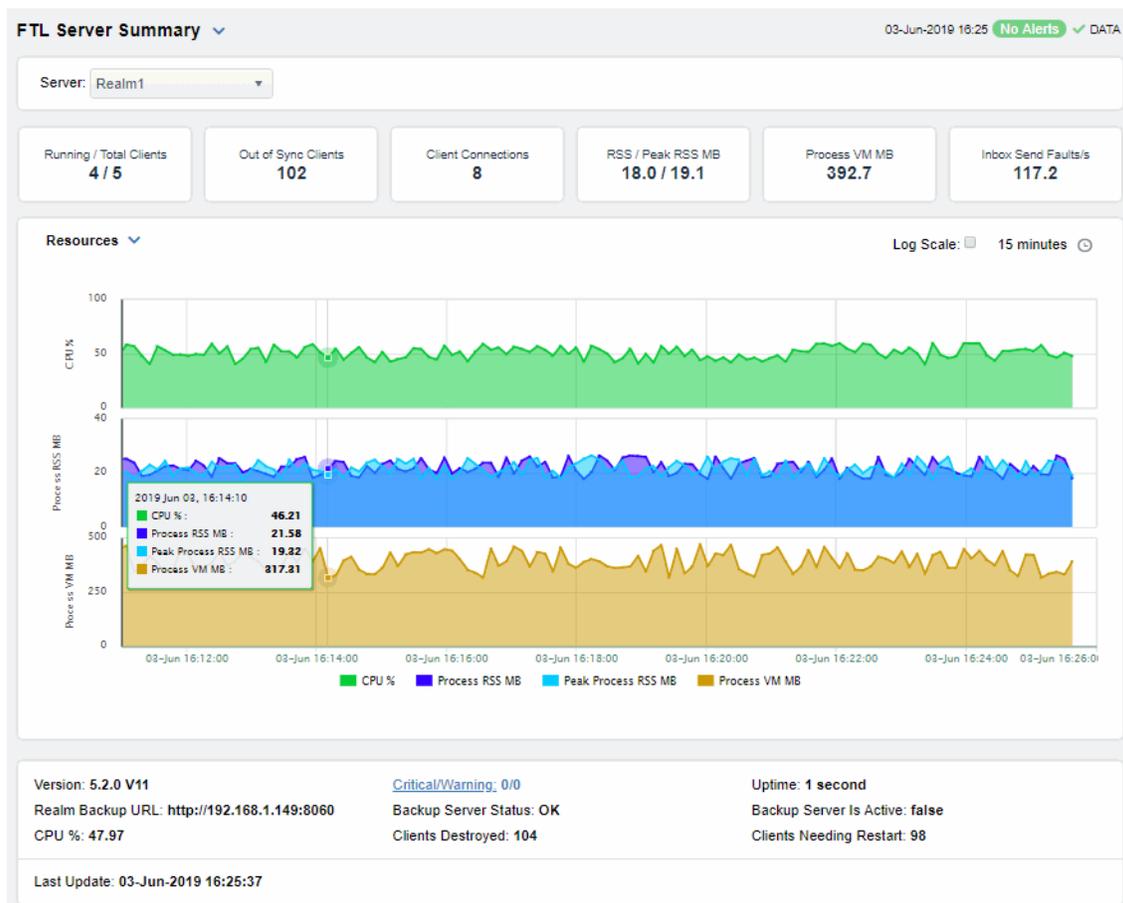
Choose from one of two trend graphs:

The **Resources** trend graph traces **CPU %** (the percent CPU used), **Process RSS** (the current amount of RSS memory used for processes), **Peak Process RSS MB** (the maximum amount of RSS memory used for processes, in megabytes) and **Process VM** (the current amount of RSS virtual memory used for processes).

The **Service Metrics** trend graph traces **Client Reconnects** (the number of reconnecting clients) and **Inbox Send Faults/s** (the number of faults when sending messages to inbox subscribers, per second\*).

You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.

Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display.



## FTL Clients - HTML

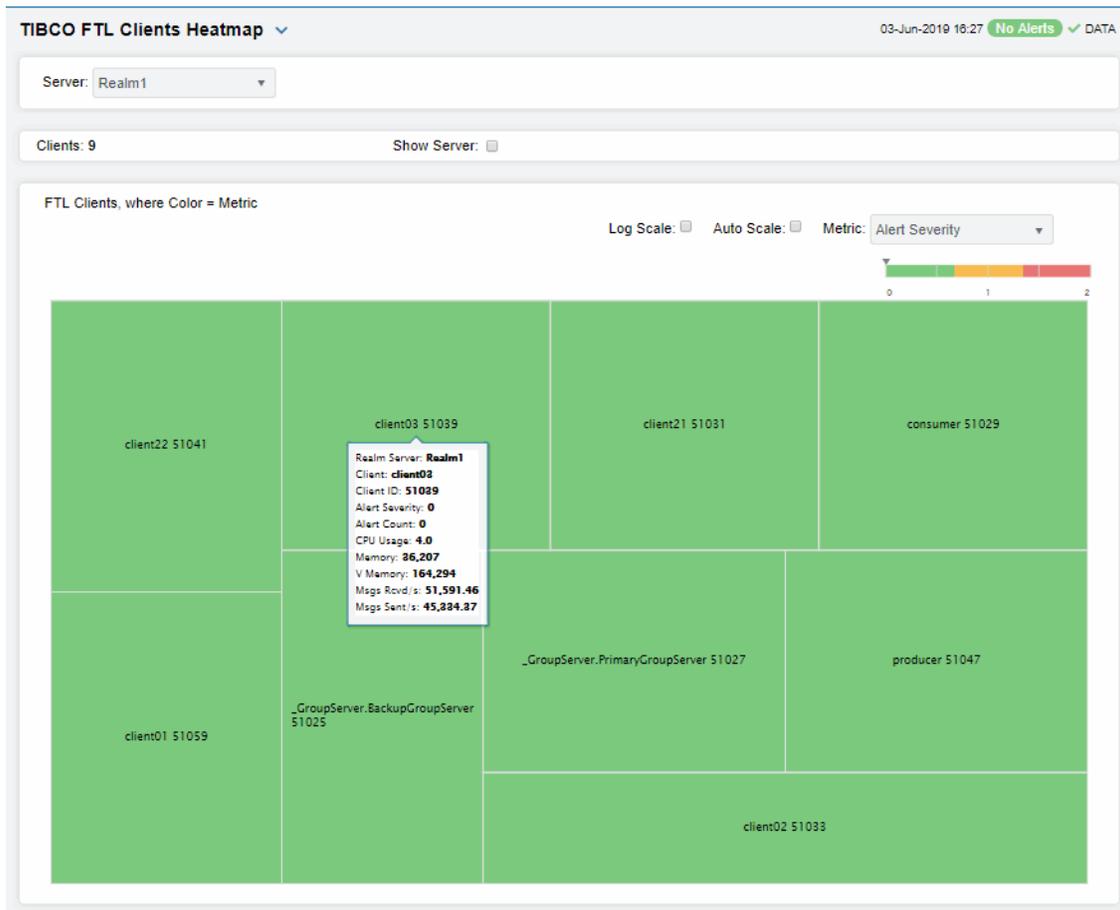
These displays present performance metrics and alert status for FTL clients. Available displays are:

- ["All Clients Heatmap - HTML"](#): Shows alert status for all FTL clients in all realms in a heatmap format.
- ["All Clients Table - HTML"](#): Shows all available utilization metrics for all FTL clients on a selected server in a tabular format.
- ["Single Client Summary - HTML"](#): Displays current and historical metrics for a single FTL client.

### All Clients Heatmap - HTML

This heatmap display allows you to quickly identify the current status of all FTL clients on all FTL servers or a particular FTL server. You can view the following metrics: the current **Alert Severity**, the current **Alert Count**, **Clients** (the total number of connected clients), **CPU Usage** (the current amount of CPU being used), **Memory** (the current amount of memory being used), **V Memory** (the current amount of virtual memory being used) and **Msgs Rcvd/s** and **Msgs Sent/s** (the number of messages received and sent per second).

Each heatmap rectangle represents a server. The rectangle color indicates the most critical alert state. You can click on a rectangle to drill-down to the ["Single Client Summary - HTML"](#) display and view metrics for that server. Clicking on the  icon next to the display title toggles between the commonly accessed **Table** and **Summary** displays. Mouse-over the rectangles to view more details about host performance and status. You can view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.

**Filter By:**

**Server** Choose the server for which you want to view data.

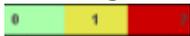
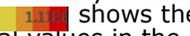
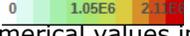
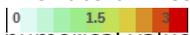
**Fields and Data**

This display includes:

**Names** Select this check box to display the names of clients.

**Log** This option enables visualization on a logarithmic scale, and should be used when the range in your data is very broad. For example, if you have data that ranges from the tens to the thousands, then data in the range of tens will be neglected visually if you do not check this option. This option makes data on both extreme ranges visible by using the logarithmic of the values rather than the actual values.

**Auto** When checked, the values of the selected metric are auto-scaled to its highest defined value. When unchecked, the values of the selected metric display based on the threshold defined for the alert associated with the selected metric. Selecting Auto helps to visualize the range of the values currently present for the selected metric instead of the threshold of the alert that has been associated with the metric. All metrics that have not been associated in the heatmap defaults with alerts use a monochromatic color gradient bar (whites and greens). All metrics that have been associated in the heatmap defaults with alerts use a multi-chromatic color gradient bar (reds, yellows, white, and greens).

<b>Metric</b>	Select the metric driving the heatmap display. The default is Alert Severity. Each <b>Metric</b> has a color gradient bar that maps values to colors. Each rectangle represents a client. Mouse-over any rectangle to display the current values of the metrics for the client. Click on a rectangle to drill-down to the associated <a href="#">"Single Client Summary - HTML"</a> display for a detailed view of metrics for that particular client.
<b>Alert Severity</b>	<p>The maximum alert level in the item (index) associated with the rectangle. Values range from <b>0</b> to <b>2</b>, as indicated in the color gradient bar , where <b>2</b> is the greatest <b>Alert Severity</b>.</p> <p><b>2</b> Metrics that have exceeded their specified <b>ALARMLEVEL</b> threshold and have an Alert Severity value of <b>2</b> are shown in red. For a given rectangle, this indicates that one or more metrics have exceeded their alarm threshold.</p> <p><b>1</b> Metrics that have exceeded their specified <b>WARNINGLEVEL</b> threshold and have an Alert Severity value of <b>1</b> are shown in yellow. For a given rectangle, this indicates that one or more metrics have exceeded their warning threshold.</p> <p><b>0</b> Metrics that have not exceeded either specified threshold have an Alert Severity value of <b>0</b> and are shown in green. For a given rectangle, this indicates that no metrics have exceeded a specified alert threshold.</p>
<b>Alert Count</b>	<p>The total number of alarm and warning alerts in a given item (index) associated with the rectangle.</p> <p>The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>CPU Usage</b>	<p>The total amount of CPU used. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the defined alert threshold of <b>TftlClientCpuUsageHigh</b>. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>Memory</b>	<p>The current amount of memory used for processing. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the defined alert threshold of <b>TftlClientMemoryHigh</b>. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>V(irtual) Memory</b>	<p>The current amount of virtual memory being used for processing. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the defined alert threshold of <b>TftlClientVirtualMemoryHigh</b>. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>Msgs Rcvd/sec</b>	<p>The rate of messages received (per second). The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the defined alert threshold of <b>TftlClientMsgsRcvdRateHigh</b>. The middle value in the gradient bar indicates the middle value of the range.</p>
<b>Msgs Sent/sec</b>	<p>The rate of messages sent (per second). The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from <b>0</b> to the defined alert threshold of <b>TftlClientMsgsSentRateHigh</b>. The middle value in the gradient bar indicates the middle value of the range.</p>

## All Clients Table - HTML

Investigate detailed utilization metrics and configuration settings for all FTL clients. The **TIBCO FTL Clients Table** contains all metrics available for clients, including the number of current client connections.

Each row in the table contains data for a particular **Realm Server**. You can search, filter, sort and choose columns to include by clicking a column header icon (to the right of each column label) and selecting **Filter**, **Sort Ascending**, **Sort Descending** or **Columns**. Or just click a column header to sort.

Right-click on a table cell to **Export to Excel** or **Copy Cell Value**.

Investigate a clients by clicking on a row to drill-down to the ["Single Client Summary - HTML"](#) display and view details for that client.

Note that the number of clients found in this table might not match the client count found in the ["FTL Servers Table - HTML"](#), possibly due to the following:

- One client might have one or more group joins resulting in a higher client count. For example, if a client has two group joins, CLIENT\_COUNT equals 3, but will only be listed as a single client in this table.
- A TIBCO bridge could have one or more logical bridges running inside the bridge process, which could result in an increased CLIENT\_COUNT even though there is actually only one client.
- Other FTL services could get a Client ID and, hence, be included in the CLIENT\_COUNT even though they are not necessarily clients.

**TIBCO FTL Clients Table** 03-Jun-2019 16:28 No Alerts DATA

Server: Realm1

Clients: 9 Filter Client Name:

Realm Server	Client	Alert Level	Alert Count	Client Status	CPU %	Process VM ...	Process RSS...	Process Pea...
Realm1	_GroupServer.BackupGroupServer	✓	0	RUNNING	5.04	167,088	27,925	29,717
Realm1	_GroupServer.PrimaryGroupServer	✓	0	RUNNING	4.83	228,259	33,858	41,089
Realm1	client01	✓	0	RUNNING	4.58	232,218	26,835	32,564
Realm1	client02	✓	0	RUNNING	4.59	213,636	34,135	39,485
Realm1	client03	✓	0	RUNNING	4.62	214,356	31,274	39,569
Realm1	client21	✓	0	RUNNING	4.97	195,559	33,503	42,511
Realm1	client22	✓	0	RUNNING	5.94	193,997	32,522	37,669
Realm1	consumer	✓	0	RUNNING	5.32	168,493	28,330	44,179
Realm1	producer	✓	0	RUNNING	4.12	182,105	36,849	35,380

## Single Client Summary - HTML

Track current and historical performance metrics for a selected FTL client. Use this display to investigate performance issues of a client.

Clicking on the metric boxes at the top of the display takes you to the ["All Clients Table - HTML"](#) display, where you can view and compare data against other clients.

Choose from one of two trend graphs:

The **Client Compute** trend graph traces **CPU %** (the percent CPU used), **Process RSS MB** (the current amount of RSS memory used for processes, in megabytes) and **Peak Process RSS MB** (the maximum amount of RSS memory used for processes, in megabytes).

The **Client Rates** trend graph traces **CPU %** (the percent CPU used), **Rcvd Msgs/s** (the number of received messages per second) and **Sent Msgs/s** ((the number of sent messages per second)).

You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.

Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display.

## TIBCO Hawk

With the Solution Package for TIBCO Hawk™ you can centralize alert events triggered by Hawk alert rule bases and performance data relevant to monitored hosts. This enables RTView to be the event correlation engine and event management system for alerts that are generated by TIBCO Hawk, RTView Enterprise and other monitoring tools such as Oracle Enterprise Manager.

The TIBCO Hawk Views displays can be found under **Components** tab > **Connectors/TIBCO Hawk**. The displays within the Views will be populated with data once the Solution Package for TIBCO Hawk is configured in the RTView DataServer for TIBCO and the RTView DataServer for TIBCO is connected to RTViewCentral. The following Views are available:

- "Hawk Agents Table"
- "Hawk Alerts Table"
- "RTView Agent Administration"

### Hawk Agents Table

This table provides a list of agents as well as network connectivity details about each agent.

Agent	Status	Last Alert Level	Cluster	IP Address	Platform
agentW46	Alive	ALERT_LOW	192.168.200.0	192.168.200.146	amd64:Windows Server 2008 R2
SLHOST93	Alive	ALERT_HIGH	192.168.200.0	192.168.200.93	amd64:Linux:2.6.32-358.11.1.el6
WIN44	Alive	ALERT_HIGH	192.168.200.0	192.168.200.144	amd64:Windows Server 2008 R2
SLHOST21(dev)	Alive	ALERT_MEDIUM	192.168.200.0	192.168.200.121	amd64:Windows 7:6.1
SLHOST5(domain5)	Alive	ALERT_MEDIUM	192.168.200.0	192.168.200.105	x86:Windows XP:5.1
SLHOST6(domain6)	Alive	ALERT_MEDIUM	192.168.200.0	192.168.200.106	x86:Windows XP:5.1
SLHOST15(sl_amx)	Alive	ALERT_HIGH	192.168.200.0	192.168.200.115	amd64:Windows 7:6.1
SLHOST17(sl_amx)	Alive	ALERT_HIGH	192.168.200.0	192.168.200.117	amd64:Windows 7:6.1
SLHOST15(sl_qa_conn)	Alive	NO_ALERT	127.0.0.0	192.168.200.115	amd64:Windows 7:6.1
SLHOST16(sl_qa_conn)	Alive	ALERT_HIGH	192.168.200.0	192.168.200.116	amd64:Windows 7:6.1
SLHOST18(sl_qa_conn)	Alive	NO_ALERT	192.168.200.0	192.168.200.118	amd64:Windows 7:6.1

#### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** **Table** open commonly accessed displays.
- 6,047** The number of items currently in the display.

- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

**Fields and Data:**

**Agent Count:** The total number of agents in the table.

**Table:**

Each row in the table is a different agent.

<b>Agent</b>	The name for the agent which is composed of the hostname and Hawk domain (in parenthesis). Agent names which do not contain an explicit Hawk domain are members of the "default" domain.
<b>Status</b>	The agent status, either <b>Alive</b> or <b>Expired</b> .
<b>Last Alert Level</b>	The most recent and most critical alert level.
<b>Cluster</b>	The IP address of the cluster to which this agent belongs.
<b>IP Address</b>	The IP subnet address for the group of machines to which this agent belongs.
<b>Platform</b>	The physical CPU class and operating system version.
<b>Last Update</b>	The date and time the row data was last updated.

**Hawk Alerts Table**

Use this display to view all Hawk alerts that have occurred in the system.

Each row in the table is a different active alert. Use the drop-down menus to filter the alerts listed. Click a column heading to sort the table on that column data.

The row color indicates the following:

**Row Color Code:**

Tables with colored rows indicate the following:

-  Red indicates that one or more alerts exceeded their ALARM LEVEL threshold in the table row.
-  Yellow indicates that one or more alerts exceeded their WARNING LEVEL threshold in the table row.
-  Green indicates that no alerts exceeded their WARNING or ALARM LEVEL threshold in the table row.

**Hawk Alerts Table** 26-Oct-2016 10:27 Data OK

Agent Filter: **All Agents**

Alert Text Filter:

Rulebase Filter:    Show Cleared Alerts **Alert Count:10**

Time	Agent	Alert ID	Alert Level	RuleBase	
26-Oct-2016 10:27:33	SLHOST6(domain6)	10	ALERT_MEDIUM	TibRV_Alerts	Received from
26-Oct-2016 10:27:19	SLHOST5(domain5)	10	ALERT_MEDIUM	TibRV_Alerts	Received from
26-Oct-2016 10:26:59	SLHOST5(domain5)	11	ALERT_LOW	System_Alerts	Server Process
26-Oct-2016 10:17:32	SLHOST17(sl_amx)	7	ALERT_HIGH	test	Current Proces
26-Oct-2016 10:06:05	SLHOST5(domain5)	13	ALERT_LOW	System_Alerts	System Uptime
26-Oct-2016 10:02:46	SLHOST16(sl_qa_conn)	4	ALERT_HIGH	generate_Alerts	Current Proces
26-Oct-2016 10:01:26	SLHOST6(domain6)	13	ALERT_LOW	System_Alerts	System Uptime
26-Oct-2016 00:26:52	SLHOST6(domain6)	11	ALERT_LOW	System_Alerts	Server Process
26-Oct-2016 00:20:33	SLHOST5(domain5)	12	ALERT_LOW	System_Alerts	Service Print S
26-Oct-2016 00:16:21	SLHOST6(domain6)	12	ALERT_LOW	System_Alerts	Service Print S

**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.
- The number of items currently in the display.

- Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

- Agent Filter** Choose one or **All Agents**.
- Alert Text Filter** Enter a string to filter alerts listed. **Clear** to remove this filter.
- Rulebase Filter** Enter a rule to filter alerts listed. **Clear** to remove this filter.
- Show Cleared Alerts** When checked, cleared alerts are included in the table.
- Alert Count** The number of alerts in the table.
- Time** The date and time the alert occurred.

- Agent** The name of the agent associated with the alert.
- Alert ID** The unique string identifier for the alert.
- Alert Level**
  - ALERT\_HIGH indicates that one or more alerts exceeded their ALARM LEVEL threshold in the table row.
  - ALERT\_MEDIUM indicates that one or more alerts exceeded their WARNING LEVEL threshold in the table row.
  - ALERT\_LOW indicates that no alerts exceeded their WARNING or ALARM LEVEL threshold in the table row.
- RuleBase** The alert system type (e.g.System\_Alerts).
- Alert Text** Descriptive text about the alert.
- Cleared** When checked, the alert has been cleared.

### RTView Agent Administration

Verify when agent metrics were last queried by the Monitor. The data in this display is predominantly used for debugging by Technical Support.

**Data Received from Remote Agents**

AgentName	AgentClass	Client ID	Total Rows Rcvd	Delta Rows rcvd	Rows Rcvd / sec	Last Receive Time
slapm	SL-RTVMGR-Agent	30002	43,412	0	0.0	10-Nov-2014 16:31:42
slapm	SL-HOSTMON-Agent	30017	53,750	35	8.6	10-Nov-2014 16:31:43
slapm	SL-BWMON-Agent	30018	423,741	8	4.0	10-Nov-2014 16:31:43
slsl4-64	SL-HOSTMON-Agent	30005	68,536	0	0.0	10-Nov-2014 16:31:37
slsl4-64	SL-BWMON-Agent	30006	91,694	0	0.0	10-Nov-2014 16:31:35
slsl4-64	SL-RTVMGR-Agent	30003	41,913	4	1.9	10-Nov-2014 16:31:43
slhost6	SL-HOSTMON-Agent	30026	23,418	0	0.0	10-Nov-2014 16:31:40
slhost6	SL-RTVMGR-Agent	30027	26,933	4	2.0	10-Nov-2014 16:31:42
slhost6	SL-BWMON-Agent	30032	26,321	14	2.3	10-Nov-2014 16:31:44
slhpux11	SL-BWMON-Agent	30012	34,363	0	0.0	10-Nov-2014 16:31:42
slhpux11	SL-HOSTMON-Agent	30010	64,394	0	0.0	10-Nov-2014 16:31:42
slhpux11	SL-RTVMGR-Agent	30011	41,820	64	15.4	10-Nov-2014 16:31:44
slvmrh2	SL-BWMON-Agent	30004	7,874	0	0.0	10-Nov-2014 16:31:38
slvmrh2	SL-RTVMGR-Agent	30001	45,352	0	0.0	10-Nov-2014 16:31:40
slvmrh2	SL-HOSTMON-Agent	30009	46,787	1	0.2	10-Nov-2014 16:31:44
slvmware	SL-BWMON-Agent	30013	6,085	0	0.0	10-Nov-2014 16:31:31
slvmware	SL-RTVMGR-Agent	30016	43,399	2	1.0	10-Nov-2014 16:31:43
slvmware	SL-HOSTMON-Agent	30015	33,434	0	0.0	10-Nov-2014 16:31:31

**Title Bar (possible features are):**

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Data Received from Remote Agents Table**

<b>AgentName</b>	Name of the agent.
<b>AgentClass</b>	Class of the agent.
<b>Client ID</b>	Unique client identifier.
<b>Total Rows Rcvd</b>	Total number of rows of data received.
<b>Rows Rcvd/sec</b>	Number of rows of data received per second.
<b>Last Receive Time</b>	Last time data was received from the agent.

---

## TIBCO Hawk - HTML

With the Solution Package for TIBCO Hawk™ - HTML you can centralize alert events triggered by Hawk alert rule bases and performance data relevant to monitored hosts. This enables RTView to be the event correlation engine and event management system for alerts that are generated by TIBCO Hawk, RTView Enterprise and other monitoring tools such as Oracle Enterprise Manager.

The TIBCO Hawk Views displays can be found under **Components** tab > **Connectors/TIBCO Hawk**. The displays within the Views will be populated with data once the Solution Package for TIBCO Hawk is configured in the RTView DataServer for TIBCO and the RTView DataServer for TIBCO is connected to RTView Central. The following displays are available:

- ["Hawk Overview - HTML"](#)
- ["Hawk Agents Table - HTML"](#)
- ["Hawk Alerts Table - HTML"](#)

### Hawk Overview - HTML

The TIBCO Hawk Overview is the top-level display for the TIBCO Hawk Connector, which provides a good starting point for immediately getting the status of all your connections on your Data Server. You can select the RTView DataServer for which you want to see data and easily view the current data for that DataServer including:

- The total number of active alerts for the selected DataServer, including the total number of critical and warning alerts.
- The number of active agents.
- A visual list of the top 10 servers containing the most total pending messages/connections/incoming messages/Async DB size in bytes on your connected DataServer.
- The total pending messages, the outgoing messages per second, and the incoming messages per second for a selected Hawk Server on your connected DataServer.

You can hover over each region in the upper half of the Overview to see more detail. You can also drill-down to see even more detail in the ["Hawk Agents Table - HTML"](#) by clicking on each respective region in the Overview.

For example, clicking on the alerts in the CRITICAL and WARNING alerts region opens the **Alerts Table by Components** display.

The bottom half of the display provides a message rates trend graph for a selected EMS server. You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.

## Hawk Agents Table - HTML

View a list of agents as well as connection details about each agent:

- **Agent:** The name for the agent which is composed of the hostname and Hawk domain (in parenthesis). Agent names which do not contain an explicit Hawk domain are members of the "default" domain.
- **Status:** The agent status, either Alive or Expired.
- **Last Alert Level:** The most recent and most critical alert level.
- **Cluster:** The IP address of the cluster to which this agent belongs.
- **IP Address:** The IP subnet address for the group of machines to which this agent belongs.
- **Platform:** The physical CPU class and operating system version.
- **Last Update:** The date and time the row data was last updated.

Each row in the table contains data for a particular agent. You can search, filter, sort and choose columns to include by clicking a column header icon (to the right of each column label) and selecting **Filter**, **Sort Ascending**, **Sort Descending** or **Columns**.

Or just click a column header to sort. Right-click on a table cell to **Export to Excel** or **Copy Cell Value**.

Hawk Agents Table 03-Jun-2019 12:16 ✓ DATA

Agents: 9

Agent	Status	Last Alert Level	Cluster	IP Address	Platform	Last Update
QAWIN5	Alive	ALERT_LOW	192.168.200.0	192.168.200.85	amd64:Windows Server 20	03-Jun-2019 11:48:02
SLHOST15(sl_qa)	Alive	ALERT_LOW	127.0.0.0	192.168.200.115	amd64:Windows 7.6.1	03-Jun-2019 12:11:25
SLHOST17(sl_qa)	Alive	ALERT_HIGH	192.168.200.0	192.168.200.117	amd64:Windows 7.6.1	03-Jun-2019 12:10:36
SLHOST18(sl_qa)	Alive	ALERT_LOW	169.254.125.0	192.168.200.118	amd64:Windows 7.6.1	03-Jun-2019 12:11:16
SLHOST5(domain)	Alive	ALERT_MEDIUM	192.168.200.0	192.168.200.105	x86:Windows XP:5.1	03-Jun-2019 12:15:42
SLHOST6(domain)	Alive	ALERT_LOW	192.168.200.0	192.168.200.106	x86:Windows XP:5.1	03-Jun-2019 11:40:24
SLHOST9	Alive	ALERT_LOW	192.168.200.0	192.168.200.109	amd64:Windows 7.6.1	03-Jun-2019 12:14:29
WIN-8-CLONE	Alive	ALERT_HIGH	192.168.200.0	192.168.200.146	amd64:Windows Server 20	03-Jun-2019 12:16:27
WIN44	Alive	ALERT_HIGH	192.168.200.0	192.168.200.144	amd64:Windows Server 20	03-Jun-2019 00:16:28

## Hawk Alerts Table - HTML

View all Hawk alerts that have occurred for a particular agent or choose **All** agents. Enter a search string in the **Alert Text Filter** or **Rulebase Filter** fields to filter alerts listed.

Each row in the table is a particular active alert.

You can select **Show Active Alerts** to remove inactive alerts from the list. You can find out the total number of alerts, the **Time** and date an alert occurred, the **Agent** name associated with the alert, the alert **ID** and the **Alert String**:

- **ALERT\_HIGH** indicates that one or more alerts exceeded their ALARM LEVEL threshold in the table row.
- **ALERT\_MEDIUM** indicates that one or more alerts exceeded their WARNING LEVEL threshold in the table row.
- **ALERT\_LOW** indicates that no alerts exceeded their WARNING or ALARM LEVEL threshold in the table row.

You can also find out the **RuleBase** for each alert (which is the alert system type, such as **System\_Alerts**), read the **Alert Text** and see if an alert is **Cleared**.

You can search, filter, sort and choose columns to include by clicking a column header icon (to the right of each column label) and selecting **Filter**, **Sort Ascending**, **Sort Descending** or **Columns**.

Or just click a column header to sort. Right-click on a table cell to **Export to Excel** or **Copy Cell Value**.

Hawk Alerts Table 03-Jun-2019 12:27 ✓ DATA [🔗](#)

Agent: - All -

Alert Text Filter:  Rulebase Filter:  Show Active Alerts Only:

Alerts: 9

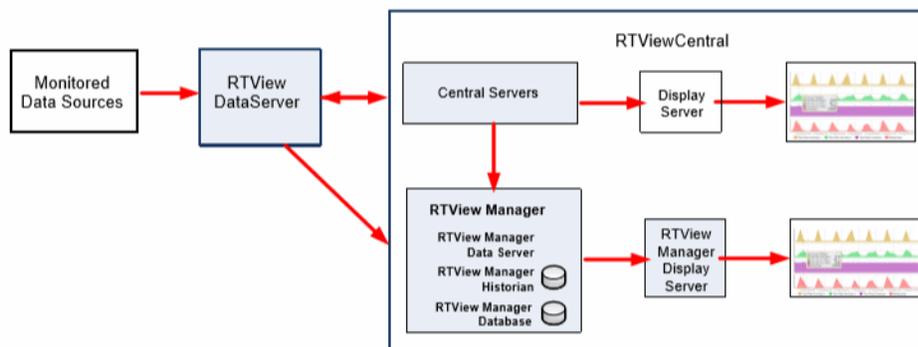
Time	Agent	Alert ID	Alert String	Alert Text	Cleared
03-Jun-2019 12:11:21	SLHOST15(sl_qa)	3	ALERT_LOW	System Uptime changed to 5 days, 4 hours, 6 minutes from last re	
03-Jun-2019 12:11:13	SLHOST18(sl_qa)	3	ALERT_LOW	Server Processes are at 66.0	
03-Jun-2019 12:17:13	SLHOST18(sl_qa)	123	ALERT_LOW	Processor load is at 83.6452697072645	✓
03-Jun-2019 12:27:06	SLHOST5(domain5)	11	ALERT_LOW	Server Processes are at 53.0	
03-Jun-2019 00:20:50	SLHOST5(domain5)	12	ALERT_LOW	Service Print Spooler is running. No Action Required.	
03-Jun-2019 12:06:20	SLHOST5(domain5)	13	ALERT_LOW	System Uptime changed to 0 days, 12 hours, 2 minutes from last r	
03-Jun-2019 12:18:15	SLHOST5(domain5)	10	ALERT_MEDIUM	Received from RV transport 7500 " tcp:7500 Advisory Message or Received from RV transport 7500 " tcp:7500 Advisory Message or	
03-Jun-2019 11:40:21	SLHOST6(domain6)	11	ALERT_LOW	System Uptime changed to 3 days, 6 hours, 1 minute from last rep	
23-Apr-2019 11:01:36	SLHOST93	3	ALERT_HIGH	no .netd are running	



## CHAPTER 9 RTView Manager

Use RTView Manager to monitor RTView Enterprise. That is, to monitor the performance of RTView processes and applications on RTViewCentral as well as any RTView DataServers to which RTViewCentral is connected.

RTView Manager runs as a separate process on RTViewCentral and has its own data server, database, Display Server and Historian.



RTView Manager requires minimal setup as required connections are auto-discovered from RTViewCentral.

There are two versions of RTView Manager, the Display Server version and the HTML version. The Display Server version is described here unless otherwise noted.

This section contains:

- ["Login to RTView Manager"](#)
- ["Overview"](#)
- ["RTView Servers Displays"](#)
- ["RTView Servers Displays - HTML"](#)
- ["JVM Displays"](#)
- ["JVM Displays - HTML"](#)
- ["Tomcat Displays"](#)
- ["Tomcat Displays - HTML"](#)

---

### Login to RTView Manager

To access this monitor, start RTViewCentral and go to **<http://localhost:3070/rview-manager-classic>**. Browse to one of the following and login (use rtvadmin/rtvadmin):

**http://<ip\_address>:3070/rtview-manager-classic** if you are running the monitor remotely.

**http://localhost:3070/rtview-manager-classic** if you are running the monitor locally.

---

## Overview

RTView Manager displays are organized by the following Views:

- **"RTView Servers Displays"**: This series of displays is for monitoring the health of the RTView servers monitoring your system. RTView Manager metrics include connected state, number of clients and other status information for Data Server, Historian and Display Server processes. Also see the **"RTView Servers Displays - HTML"** version.
- **"JVM Displays"**: For monitoring the health of Java Virtual Machine (JVM) processes. JVM metrics track garbage collection information and trends, including memory usage before and after garbage collection, duration and duty cycles. This, combined with tracking of JVM memory pool trends, enables you to be notified of memory leaks, unusual garbage collection activities and CPU and memory resource issues automatically with minimal user analysis, speeding the discovery of the root cause of any issue. It also monitors a Java Virtual Machine's memory heap, non-heap memory, monitor threads and other key metrics to ensure the JVM has good performance. Detailed metrics including JVM CPU usage, Max Heap, Current Heap, Used Heap and Live Threads can all be tracked over time. Also see the **"JVM Displays - HTML"** version.
- **"Tomcat Displays"**: For monitoring the health of Tomcat servers, applications and all installed web modules. Performance data provided includes current and historic metrics, number of sessions, request rates, cache hit rates and data transmission metrics. Also see the **"Tomcat Displays - HTML"** version.

---

## RTView Servers Displays

The following RTView Servers displays can be found under **Components** tab > **Processes** > **RTView Processes** > **RTView Servers** after installation.

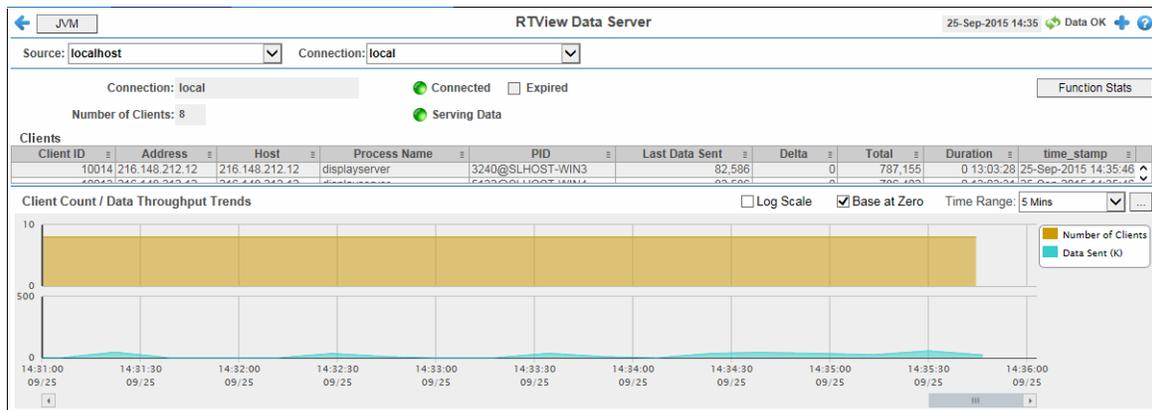
These displays present performance data for all RTView Enterprise servers. Use these displays to monitor the health of the servers monitoring your system. Displays in this View are:

- **"Data Servers"**: Shows metrics for RTView Data Servers.
- **"Display Servers"**: Shows metrics for RTView Display Servers.
- **"Historian Servers"**: Shows metrics for RTView Historian Servers.
- **"Version Info"**: Shows the version information of each jar used in each connected RTView application.

### Data Servers

Track data transfer metrics for RTView Data Servers, client count and throughput trends.

Use the available drop-down menus or right-click to filter data shown in the display.



#### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

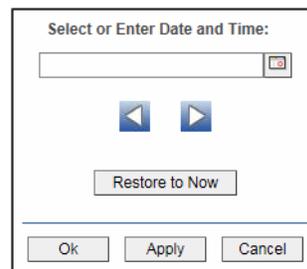
- Source** Select the type of connection to the RTView Server.
- Connection** Select an RTView Server from the drop-down menu. Names can be modified in the RTView Server configuration properties file.
- Connection** The connection selected from the **Connection** drop-down menu.
- Number of Clients** The number of clients currently server on this Data Server.
- Connected** The Data Server connection state:
  - Disconnected.
  - Connected.
- Serving Data**
  - The Data Server is not currently serving data.
  - The Data Server is currently serving data.
- Expired** This server has been marked as expired after no activity.
- Function Stats** Opens the **RTView Function Stats** display which shows detailed performance statistics for RTView functions in the selected Data Server. This button is only enabled if the RTView has a JMX connection defined for the selected Data Server.
- Clients** This table describes all clients on the selected server.

<b>Address</b>	The client IP address.
<b>Client ID</b>	The unique client identifier.
<b>Duration</b>	The amount of time for this client session. Format: <b>dd HH:MM:SS</b> <b>&lt;days&gt; &lt;hours&gt;:&lt;minutes&gt;:&lt;seconds&gt;</b> <b>For example:</b> <b>10d 08:41:38</b>
<b>Host</b>	The client host name.
<b>Last Data Sent</b>	The amount of data, in bytes, last sent to the client.
<b>Delta</b>	The amount of data, in bytes, sent since the last update.
<b>Total</b>	The total amount of data, in bytes, sent to the client.
<b>TIME_STAMP</b>	The date and time this row of data was last updated.

### Client Count / Data Throughput Trends

Shows throughput metrics for all clients on the selected server.

<b>Log Scale</b>	Enable to use a logarithmic scale for the Y axis. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.
<b>Base at Zero</b>	Use zero as the Y axis minimum for all graph traces.
<b>Time Range</b>	Select a time range from the drop down menu varying from <b>2 Minutes</b> to <b>Last 7 Days</b> , or display <b>All Data</b> . To specify a time range, click Calendar <input type="button" value="..."/> .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

<b>Number of Clients</b>	Traces the number of clients being served by the Data Server.
<b>Data Sent</b>	Traces the total amount of data, in Kilobytes, sent to all clients.

## Display Servers

Track display utilization metrics for RTView Display Servers.

Use the available drop-down menus or right-click to filter data shown in the display.

The screenshot shows the RTView Display Server interface. At the top, it displays 'JVM' and 'RTView Display Server' with a timestamp of '07-Mar-2019 15:16' and a 'Data OK' status. Below this, there are three dropdown menus: 'Source' set to 'localhost', 'Connection' set to 'RTVMGR\_DISPLAYSERVER', and 'DataServer'. A status bar indicates 'Connected' with a green dot and 'Expired' with a grey dot, along with a 'Function Stats' button. Configuration fields include 'Display Timeout (seconds): 60', 'Image Quality (0 - 100): 75', 'Number of Active Displays: 7', 'Maximum Number of Active Displays: 80', and 'Sessions with Active Displays: 3'. A table titled 'Display Data' shows a list of active displays with columns for Display Name, Session, Panel ID, Expired, and SUI. The table contains several rows, with the first row highlighted in blue.

Display Name	Session	Panel ID	Expired	SUI
rtv_about.rtv	88cf521fa8549f3aa8900bd800256b8	e8d5647b33464822bf8c5cct0f9224f9	<input type="checkbox"/>	SrvPopFlag 0 SrvLast
rtvmgr_title.rtv	88cf521fa8549f3aa8900bd800256b8	3e28aDe2e3de4430afb6ab3fb8f30caa	<input type="checkbox"/>	ScurrentDisplay.rtvmgr_
rtvmgr_title.rtv	6faf757945ba4d719f0d804acea856d1	69a4f7c6edee4b99bea6b2e09531b58	<input type="checkbox"/>	ScurrentDisplay.rtvmgr_
rtvmgr_navtree.rtv	88cf521fa8549f3aa8900bd800256b8	9bb8a1247ddb4f0945ed3aace4eb1f3	<input type="checkbox"/>	ScurrentDisplay.rtvmgr_
rtvmgr_navtree.rtv	3cb4973f040c4699aee93a5c10308c30	c96ebf0f53d341aeb9d22ce2afcb01	<input type="checkbox"/>	ScurrentDisplay.rtvmgr_
jvm_mempool_trends.rtv	88cf521fa8549f3aa8900bd800256b8	a6b9d29affde471ea2d3ded48d0ce9e5	<input type="checkbox"/>	ScolorColumnNameAct
rtv_server_summary_display.rtv	88cf521fa8549f3aa8900bd800256b8	bce9973b5cda4890acea9b2be7684d13	<input type="checkbox"/>	ScolName:" SvalueFact

### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

### Fields and Data

This display includes:

- Source** Select the type of connection to the RTView Server.
- Connection** Select an RTView Server from the drop-down menu. Names can be modified in the RTView Server configuration properties file.
- Connected** The Display Server connection state:
  - Disconnected.
  - Connected.
- Expired** This server has been marked as expired after no activity.
- Function Stats** Opens the **RTView Function Stats** display which shows detailed performance statistics for RTView functions in the selected Display Server. This button is only enabled if the RTVMGR has a JMX connection defined for the selected Display Server.

<b>Display Timeout (seconds)</b>	The amount of time, in seconds, that a display can be kept in memory after the Display Servlet has stopped requesting it. The default is <b>60</b> seconds (to allow faster load time when switching between displays).
<b>Image Quality (0-100)</b>	A value between <b>0</b> and <b>100</b> , which controls the quality of the generated images. If the value is <b>100</b> , the Display Server outputs the highest quality image with the lowest compression. If the value is <b>0</b> , the Display Server outputs the lowest quality image using the highest compression. The default is <b>75</b> .
<b>Number of Active Displays</b>	The total number of displays currently being viewed by a user.
<b>Maximum Number of Active Displays</b>	The maximum number of displays kept in memory. The default is <b>20</b> (to optimize memory used by the Display Server).
<b>Sessions with Active Displays</b>	Number of clients accessing the Display Server.

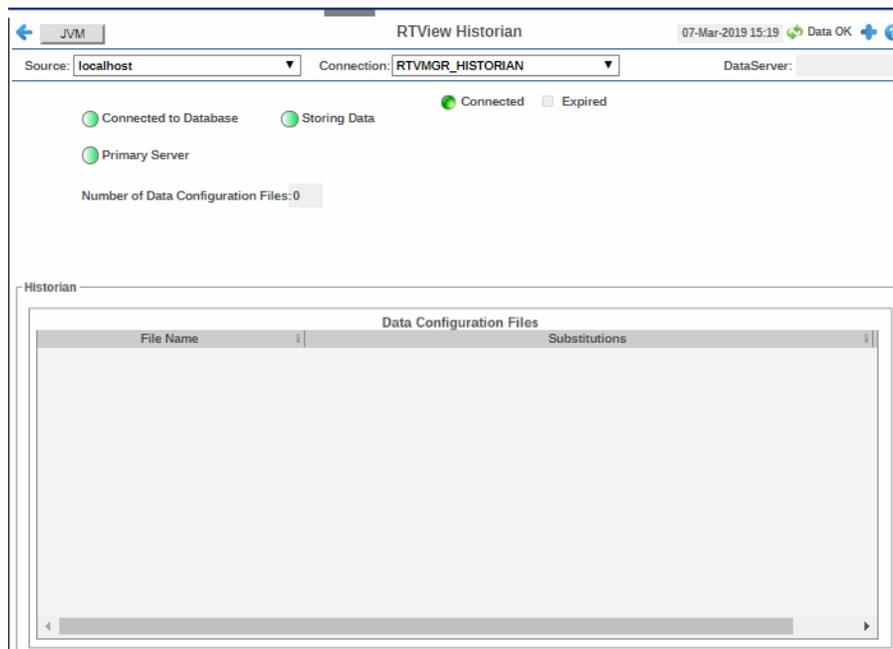
#### Display Data / Active Displays

<b>Display Name</b>	The name of the currently open display.
<b>Session</b>	A unique string identifier assigned to each session.
<b>Panel ID</b>	A unique string identifier assigned to each panel. The Display Server loads each display requested by each client into a panel. This ID can be useful in troubleshooting.
<b>Substitutions</b>	Lists the substitutions used for the display.
<b>Last Ref</b>	The amount of time that has elapsed since the display was last requested by a client.
<b>ID</b>	The client ID.
<b>Preloaded</b>	When checked, indicates that the display ( <b>.rtv</b> ) file is configured in the <b>DISPLAYSERVER.ini</b> file to be preloaded. The <b>history_config</b> option is used to configure display preloading. Preloading a display makes data immediately available. Preloaded displays are not unloaded unless the Display Server is restarted or the display cache is cleared via JMX. This option can be used multiple times to specify multiple displays to preload.

## Historian Servers

Track the status of RTView Historian Servers and data configuration file usage. View the caches that are archived by the Historian application, substitution variables associated with the history cache configuration file, as well as the history cache status. You can also stop and start the Historian, and purge data.

Use the available drop-down menus or right-click to filter data shown in the display.



#### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

### Fields and Data

This display includes:

- Source** Select the type of connection to the RTView Server.
- Connection** Select an RTView Server from the drop-down menu. Names can be modified in the RTView Server configuration properties file.
- Connected** The Historian Server connection state:
  - Disconnected.
  - Connected.
- Expired** This server has been marked as expired after no activity.
- Connected to Database** The Historian Server database connection state:
  - Disconnected.
  - Connected.

**Primary Server**

When green, indicates that this Historian, when used within a group of Historians, is the primary group member. If the primary member fails or shuts down, the standby member with the highest priority becomes the primary group member. When red, indicates that the Historian is a secondary server.

The Historian Server member state:

- The Historian Server is a secondary group member.
- This Historian is the primary group member.

**Number of Data Configuration Files**

The number of configuration files that are used by the history cache.

**Historian / Data Configuration Files**

**File Name** The name of the history cache configuration file.

**Substitutions** Lists the substitutions specified in the history cache configuration file.

**Version Info**

This display provides detailed version information for all of the connected RTView applications. You can view specific applications by filtering data using the **Source**, **Connection**, **Filter Field**, and **Filter Value** fields at the top of the display. This display provides valuable information about the version of each jar that is used in each connected RTView application that can be used to help Technical Support when issues arise. All RTView applications use multiple jars and this display lists the version information for each jar in the application. The **ApplicationConfiguration** column shows the version of the jar that contains the main class for the application which is also the version that is printed to the console at startup. The **JarConfiguration** shows the version of the jar specified in the **JarName** field. When **ApplicationConfiguration** and **JarConfiguration** do not match, it indicates that the application is using jars from multiple releases of RTView or that the application is using a patched jar. Rows in the table where the **JarConfiguration** does not match the **ApplicationConfiguration** are highlighted in teal.

**Note:** RTView applications running versions previous to this enhancement will only have one row in the table and will display "version info not supported in this version" in the **ApplicationConfiguration** column.

Detailed Version for All Connected RTView Applications  
Rows where the JarConfiguration does not match ApplicationConfiguration are highlighted in teal

Source	Connection	ApplicationName	JarName	ApplicationConfiguration	JarConfiguration	JarVersionNumb
WIN3	SLMON-DISP-5	RTView Display Server	gmsjagentds.jar	APM.3.0.0.0_20150910_000.19559-alpha_119	APM.3.0.0.0_20150910_000.19559-alpha_119	3.0.0.0
WIN3	SLMON-DISP-5	RTView Display Server	gmsjalertds.jar	APM.3.0.0.0_20150910_000.19559-alpha_119	APM.3.0.0.0_20150910_000.19559-alpha_119	3.0.0.0
WIN3	SLMON-DISP-5	RTView Display Server	gmsjcacheds.jar	APM.3.0.0.0_20150910_000.19559-alpha_119	APM.3.0.0.0_20150910_000.19559-alpha_119	3.0.0.0
WIN3	SLMON-DISP-5	RTView Display Server	gmsjcmdbds.jar	APM.3.0.0.0_20150910_000.19559-alpha_119	APM.3.0.0.0_20150910_000.19559-alpha_119	3.0.0.0
WIN3	SLMON-DISP-5	RTView Display Server	gmsjext.jar	APM.3.0.0.0_20150910_000.19559-alpha_119	APM.3.0.0.0_20150910_000.19559-alpha_119	3.0.0.0
WIN3	SLMON-DISP-5	RTView Display Server	gmsjflash.jar	APM.3.0.0.0_20150910_000.19559-alpha_119	APM.3.0.0.0_20150910_000.19559-alpha_119	3.0.0.0
WIN3	SLMON-DISP-5	RTView Display Server	gmsjlog4jds.jar	APM.3.0.0.0_20150910_000.19559-alpha_119	APM.3.0.0.0_20150910_000.19559-alpha_119	3.0.0.0
WIN3	SLMON-DISP-5	RTView Display Server	gmsjmodels.jar	APM.3.0.0.0_20150910_000.19559-alpha_119	APM.3.0.0.0_20150910_000.19559-alpha_119	3.0.0.0
WIN3	SLMON-DISP-5	RTView Display Server	gmsjlapds.jar	APM.3.0.0.0_20150910_000.19559-alpha_119	APM.3.0.0.0_20150910_000.19559-alpha_119	3.0.0.0
WIN3	SLMON-DISP-5	RTView Display Server	gmsjpipeps.jar	APM.3.0.0.0_20150910_000.19559-alpha_119	APM.3.0.0.0_20150910_000.19559-alpha_119	3.0.0.0
WIN3	SLMON-DISP-5	RTView Display Server	gmsjrrds.jar	APM.3.0.0.0_20150910_000.19559-alpha_119	APM.3.0.0.0_20150910_000.19559-alpha_119	3.0.0.0
WIN3	SLMON-DISP-5	RTView Display Server	gmsjrtvhistorian.jar	APM.3.0.0.0_20150910_000.19559-alpha_119	APM.3.0.0.0_20150910_000.19559-alpha_119	3.0.0.0
WIN3	SLMON-DISP-5	RTView Display Server	gmsjrtvviewer.jar	APM.3.0.0.0_20150910_000.19559-alpha_119	APM.3.0.0.0_20150910_000.19559-alpha_119	3.0.0.0



## Fields and Data

This display includes:

<b>Source</b>	Select a filter value for the <b>Source</b> column.
<b>Connection</b>	Select a filter value for the <b>Connection</b> column.
<b>Filter Field</b>	Select a table column from the drop-down menu to perform a search in: <b>ApplicationName, JarName, ApplicationConfiguration, JarConfiguration, JarVersionNumber, JarVersionDate, JarReleaseDate, and JarMicroVersion.</b> Filters limit display content and drop-down menu selections to only those items that pass through the selected filter's criteria. If no items match the filter, you might have zero search results (an empty table). Double-clicking on a specific field in the table will populate this field with the selected field's content. For example, double-clicking on the <b>DataServerName</b> field in one of the rows displays the entire field's content into this field.
<b>Clear</b>	Clears entries in the <b>Filter Field</b> display list, <b>Filter Value</b> field, and <b>Not Equal</b> check box.
<b>Filter Value</b>	Enter the (case-sensitive) string to search for in the selected <b>Filter Field</b> .
<b>RegEx</b>	Select to use the <b>Filter Value</b> as a regular expression when filtering. When selected, the <b>Not Equal</b> check box displays.
<b>Not Equal</b>	Works in conjunction with the <b>RegEx</b> field. Selecting this check box searches for values in the specified <b>Filter Field</b> that are NOT equal to the value defined in the <b>Filter Value</b> field. For example, if the <b>Filter Field</b> specified is <b>JarMicroVersion</b> , the <b>Filter Value</b> is specified as <b>317</b> , and this check box is selected, then only those rows containing <b>JarMicroVersion</b> fields NOT EQUAL to <b>317</b> will display. This field is only enabled when the <b>RegEx</b> check box is checked.
<b>Source</b>	The name of the source of the RTView .
<b>Connection</b>	Lists the name of the JMX connection to the RTView application.
<b>Application Name</b>	Lists the name of the application.
<b>JarName</b>	Lists the name of the jar used in the connected application.
<b>Application Configuration</b>	Lists the configuration string of the application. This string contains the main application version that corresponds to the version information printed to the console at startup.
<b>JarConfiguration</b>	Lists the configuration string for the jar.
<b>JarVersionNumber</b>	Lists the version number for the jar.
<b>JarVersionDate</b>	Lists the version date for the jar.
<b>JarReleaseType</b>	Lists the release type for the jar.
<b>JarMicroVersion</b>	Lists the micro version for the jar.

<b>Expired</b>	When checked, this connection is expired due to inactivity.
<b>time_stamp</b>	The time at which the information in the current row was last received.
<b>DataServerName</b>	The name of the RTView Data Server connection.

---

## RTView Servers Displays - HTML

The following RTView Servers displays can be found under **Components** tab > **Processes** > **JVM Processes** > **RTView Servers** after installation.

These displays present performance data for all RTView servers. Use these displays to monitor the health of the servers monitoring your system. Displays are:

- ["Data Servers - HTML"](#): Shows metrics for RTView Data Servers.
- ["Data Server Summary - HTML"](#): Shows details for one Data Server.
- ["Display Servers - HTML"](#): Shows metrics for RTView Display Servers.
- ["Display Server Summary - HTML"](#): Shows details for one Display Server.
- ["Historian Servers - HTML"](#): Shows metrics for RTView Historian Servers.

### Data Servers - HTML

View connections on one or all RTView Data Servers, as well as connection status and client count. Choose one or **All** data servers from the **Source:** drop-down menu. Each row in the table is a different data server. You can search, filter, sort and choose columns to include by clicking a column header icon (to the right of each column label) and selecting Filter, Sort Ascending, Sort Descending or Columns.

Or just click a column header to sort. Right-click on a table cell to Export to Excel or Copy Cell Value.

Double-click to drill-down to details about the selected data server as well as the selected connection in the **Data Server Summary** display.

The screenshot shows a web interface titled "RTView Data Servers Table" with a timestamp of "16-Apr-2019 14:56" and a "DATA" indicator. Below the title is a "Source" dropdown menu set to "- All -". A status bar indicates "Data Servers: 1". The main content is a table with the following data:

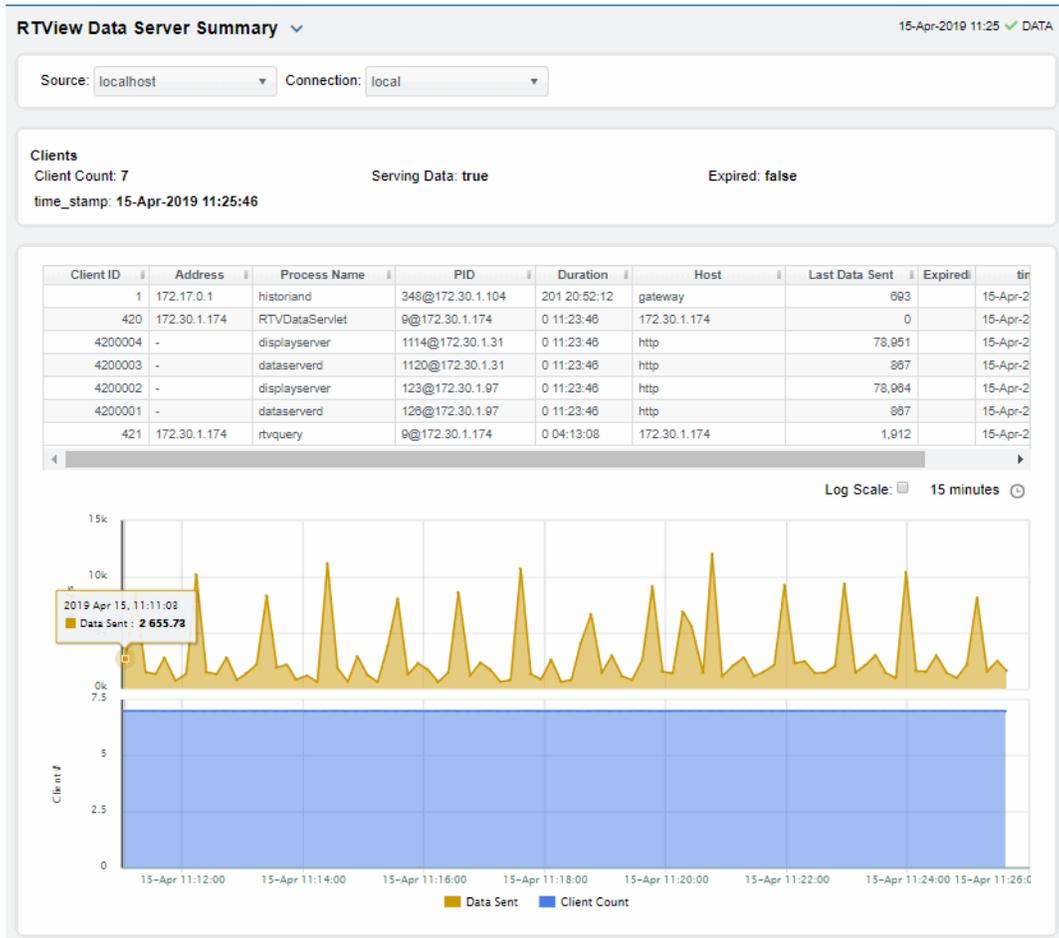
Source	Connection	Client Count	Serving Data	Expired	time_stamp
localhost	local	7	✓		16-Apr-2019 14:56:16

## Data Server Summary - HTML

Track utilization metrics for a specific data server and a connection. Choose a data server and a connection from the **Source** and **Connection** drop-down menus. View client details such as client ID, IP address, process name, host and duration. The trend graph traces data sent and client count for the selected connection. You can hover over the trend graph to see the values at a particular time. You can specify the

time range for the trend graph and view data based on a log scale, which enables visualization

on a logarithmic scale and should be used when the range in your data is very broad.



## Display Servers - HTML

Track display utilization metrics for connections on your RTView Display Servers. Each row in the table is for a particular connection. Select a **Source** or **All** sources to see connection values for:

- **Display Timeout (seconds)** The amount of time, in seconds, that a display can be kept in memory after the Display Servlet has stopped requesting it. The default is **60** seconds (to allow faster load time when switching between displays).
- **Image Quality (0-100)** A value between **0** and **100**, which controls the quality of the generated images. If the value is **100**, the Display Server outputs the highest quality image with the lowest compression. If the value is **0**, the Display Server outputs the lowest quality image using the highest compression. The default is **75**.
- **Number of Active Displays** The total number of displays currently being viewed by a user.
- **Maximum Number of Active Displays** The maximum number of displays kept in memory. The default is **20** (to optimize memory used by the Display Server).
- **Sessions with Active Displays** Number of clients accessing the Display Server.
- **Expired** This server has been marked as expired after no activity.

Investigate by clicking a row to see details for that connection in the summary display.

RTView Display Servers Table 16-May-2019 14:57 ✓ DATA

Source: localhost

Display Servers: 3

Source	Connection	Display Timeout	Image Quality	Max Displays	Display Count	Expired	time_stamp
localhost	EM-DISP-X	60	75	80	0		16-May-2019 14:57:00
localhost	EM-DISP-X2	60	75	80	0		16-May-2019 14:57:00
localhost	EMSMON-DISP-3	60	75	80	0		16-May-2019 14:57:00

## Display Server Summary - HTML

Track display utilization for connections on a particular RTView Display Server. Each row in the table is a different connection. Select a **Source** and a connection to see values for:

- **Display Name** The name of the currently open display.
- **Session** A unique string identifier assigned to each session.
- **Panel ID** A unique string identifier assigned to each panel. The Display Server loads each display requested by each client into a panel. This ID can be useful in troubleshooting.
- **Substitutions** Lists the substitutions used for the display.
- **Last Ref** The amount of time that has elapsed since the display was last requested by a client.
- **ID** The client ID.
- **Preloaded** When checked, indicates that the display (.rtv) file is configured in the **DISPLAYSERVER.ini** file to be preloaded. The `history_config` option is used to configure display preloading. Preloading a display makes data immediately available. Preloaded displays are not unloaded unless the Display Server is restarted or the display cache is cleared via JMX. This option can be used multiple times to specify multiple displays to preload.

RTView Display Servers Table 16-May-2019 14:57 ✓ DATA

Source: localhost

Display Servers: 3

Source	Connection	Display Timeout	Image Quality	Max Displays	Display Count	Expired	time_stamp
localhost	EM-DISP-X	60	75	80	0		16-May-2019 14:57:00
localhost	EM-DISP-X2	60	75	80	0		16-May-2019 14:57:00
localhost	EMSMON-DISP-3	60	75	80	0		16-May-2019 14:57:00

## Historian Servers - HTML

Track the status of RTView Historian Servers, their connections, status and role and data configuration file usage. View the caches that are archived by the Historian application, substitution variables associated with the history cache configuration file, as well as the history cache status.

Each row in the table contains data for a particular server. You can search, filter, sort and choose columns to include by clicking a column header icon (to the right of each column label) and selecting Filter, Sort Ascending, Sort Descending or Columns. Or just click a column header to sort. Right-click on a table cell to Export to Excel or Copy Cell Value.

RTView Historians Table 16-May-2019 15:12 ✓ DATA

Source: localhost

Historians: 2

Source	Connection	Connected To DB	Suspended	Config File Count	Primary	Expired	Time Stamp
localhost	RTVMGR-HIST-X2			0	✓		16-May-2019 15:12:26
localhost	EMSMON-HIST-3			0	✓		16-May-2019 15:12:26

## JVM Displays

The RTView Manager JVM displays present performance data for monitored Java Virtual Machine (JVM) processes. Use these displays to examine the current and historical performance metrics and resource usage of JVMs. Any JVM that is enabled for monitoring can be included in these displays. The displays include summary overviews and detail pages with historical trends.

You can set alert thresholds on performance and resource metrics for your JVMs, including **CPU Percent**, **Memory Used** and **Gc Duty cycle**. Alerts are shown in the ["All JVMs Heatmap"](#) display. Use the detailed JVM displays to investigate further; for example a **Memory Used** alarm might take you to the ["JVM Summary"](#) display to get historical memory use, or a **Gc Duty Cycle** alarm might take you to the ["JVM GC Trends"](#) display. Displays in this View are:

The following JVM Views can be found under **Components** tab > **Processes /JVM Processes > JVM** once RTView Manager is installed:

- ["All JVMs"](#)
- ["Single JVM"](#)



**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.

open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Fields and Data**

This display includes:

**Source** Choose one or **All Sources** to display.

**JVM Count** The number of JVM connections shown in the display.

**Show Inactive** Select to show inactive connections.

**Connection** Select to show JVM connections names.

**Metric**

Select the Metric to display in the heatmap. Each Metric has a color gradient bar that maps relative values to colors.

- Alert Severity** The maximum level of alerts in the heatmap rectangle. Values range from **0** - **2**, as indicated in the color gradient bar, where **2** is the highest Alert Severity.

  - Red indicates that one or more alerts have reached their alarm threshold. Alerts that have exceeded their specified ALARM LEVEL threshold have an Alert Severity value of **2**.
  - Yellow indicates that one or more alerts have reached their alarm threshold. Alerts that have exceeded their specified WARNING LEVEL threshold have an Alert Severity value of **1**.
  - Green indicates that no alerts have reached their alert thresholds. Alerts that have not exceeded their specified thresholds have an Alert Severity value of **0**.
- Alert Count** The number of alerts for the rectangle. The color gradient bar values range from **0** to the maximum number of alerts in the heatmap.
- CPU %** The total percent (%) CPU utilization for the rectangle. The color gradient bar values range from **0** to the maximum percent (%) CPU utilization in the heatmap.
- Memory %** The total percent (%) memory utilization for the rectangle. The color gradient bar values range from **0** to the maximum percent (%) memory utilization in the heatmap.
- Current Heap** The current amount of heap committed for the connection, in kilobytes. The color gradient bar values range from **0** to the maximum amount in the heatmap.
- Used Heap** The total amount of heap used by the connection, in kilobytes. The color gradient bar values range from **0** to the maximum amount used in the heatmap.

**All JVMs Table**

View JVM connection details for one or all sources, the most critical alert state for each JVM connection, as well as CPU and memory utilization in a tabular format. Each row in the table is a different connection.

Choose one or **All Sources** from the **Source** drop-down menu. Check the **Show Inactive** box to include inactive connections. The row color for inactive connections is dark red. Click Sort  to order column data. Drill-down and investigate by clicking a row in the table to view details for the selected connection in the **JVM Summary** display.

← Heatmap
All JVMs - Table View
19-Jan-2017 14:01  Data OK 

Source: All Sources

JVM Count: 56  Show Inactive

**All JMX Connections**

Connection	Source	Expired	Connected	Alert Severity	Alert Count	Host	Port	
ALERT_SERVER	localhost	<input type="checkbox"/>		0	0	localhost	10023	102
ALERT_SERVER	TBSender	<input type="checkbox"/>		0	0	localhost	10023	102
ALERTHISTORIAN	localhost	<input type="checkbox"/>		0	0	localhost	10025	110
ALERTHISTORIAN	TBSender	<input type="checkbox"/>		0	0	localhost	10025	110
AMXMON-alpha-TB34	localhost	<input type="checkbox"/>		0	0	192.168.200.34	6368	309
AMXMON-alpha-TB34	TBSender	<input type="checkbox"/>		0	0	192.168.200.34	6368	309
AMXMON-alpha-TB34-HIST	localhost	<input type="checkbox"/>		0	0	192.168.200.34	6367	633
AMXMON-beta-TB3-HIST	localhost	<input type="checkbox"/>		0	0	192.168.200.133	6367	473
BWMON-alpha-TB34	localhost	<input type="checkbox"/>		0	0	192.168.200.34	3368	321
BWMON-alpha-TB34	TBSender	<input type="checkbox"/>		0	0	192.168.200.34	3368	321
BWMON-alpha-TB34-HIST	localhost	<input type="checkbox"/>		0	0	192.168.200.34	3367	325
BWMONITOR-release-WIN-8	localhost	<input type="checkbox"/>		0	0	192.168.200.146	3368	904
BWMONITOR-TB8	localhost	<input type="checkbox"/>		0	0	192.168.200.138	3368	270
CONFIG_SERVER	localhost	<input type="checkbox"/>		0	0	localhost	10013	990
CONFIG_SERVER	TBSender	<input type="checkbox"/>		0	0	localhost	10013	990
DISPLAYSERVER	localhost	<input type="checkbox"/>		0	0	localhost	10024	100
DISPLAYSERVER	TBSender	<input type="checkbox"/>		0	0	localhost	10024	100
DISPLAYSERVER_DARKSTY	localhost	<input type="checkbox"/>		0	0	localhost	10124	118
DISPLAYSERVER_DARKSTY	TBSender	<input type="checkbox"/>		0	0	localhost	10124	118
EMSMON_SENDER-alpha-TB	TBSender	<input type="checkbox"/>		0	0	192.168.200.34	3166	289
EMSMON_SENDER-alpha-TB	localhost	<input type="checkbox"/>		0	0	192.168.200.34	3166	289

#### Title Bar (possible features are):

-   Open the previous and upper display.
-  Open an instance of this display in a new window.
-  Open the online help page for this display.
- Menu Table open commonly accessed displays.

 Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

#### Row Color Code:

Tables with colored rows indicate the following:

-  Red indicates that one or more alerts exceeded their ALARM LEVEL threshold in the table row.
-  Yellow indicates that one or more alerts exceeded their WARNING LEVEL threshold in the table row.
-  Green indicates that no alerts exceeded their WARNING or ALARM LEVEL threshold in the table row.

#### Fields and Data

This display includes:

- Source** Choose one or **All Sources** to display.

**JVM Count:** The number of JVM connections in the table.

**Show Inactive** Select to include inactive connections.

#### All JMX Connections Table

<b>Connection</b>	The name of the JVM connection.
<b>Source</b>	The name of the source.
<b>Expired</b>	When checked, this connection is expired due to inactivity.
<b>Connected</b>	The data connection state:  Disconnected.  Connected.
<b>Alert Severity</b>	The maximum level of alerts associated with the connection. Values range from <b>0</b> to <b>2</b> , where <b>2</b> is the greatest Alert Severity.  One or more alerts associated with the connection exceeded their ALARM LEVEL threshold.  One or more alerts associated with the connection exceeded their WARNING LEVEL threshold.  No alerts associated with the connection have exceeded their thresholds.
<b>Alert Count</b>	The current number of alerts for the connection.
<b>Host</b>	The name of the host for this connection.
<b>Port</b>	The port number for the connection.
<b>PID</b>	The connection PID.
<b>CPU %</b>	The amount of CPU, in percent (%) used by this connection.
<b>Max Heap</b>	The maximum amount of heap used by this connection, in kilobytes.
<b>Current Heap</b>	The current amount of committed heap for this connection, in kilobytes.
<b>Used Heap</b>	The current amount of heap used by this connection, in kilobytes.
<b>Mem % Used</b>	The amount of JVM memory used by this connection, in percent (%).
<b>RtvAppType</b>	The type of RTView application, where: <b>1</b> is for the Historian, <b>3</b> is for the Data Server; <b>5</b> is for the Display Server, and <b>0</b> is a non-RTView application.
<b>Source</b>	The Data Server that sent this value.
<b>time_stamp</b>	The date and time this row of data was last updated.

## Single JVM

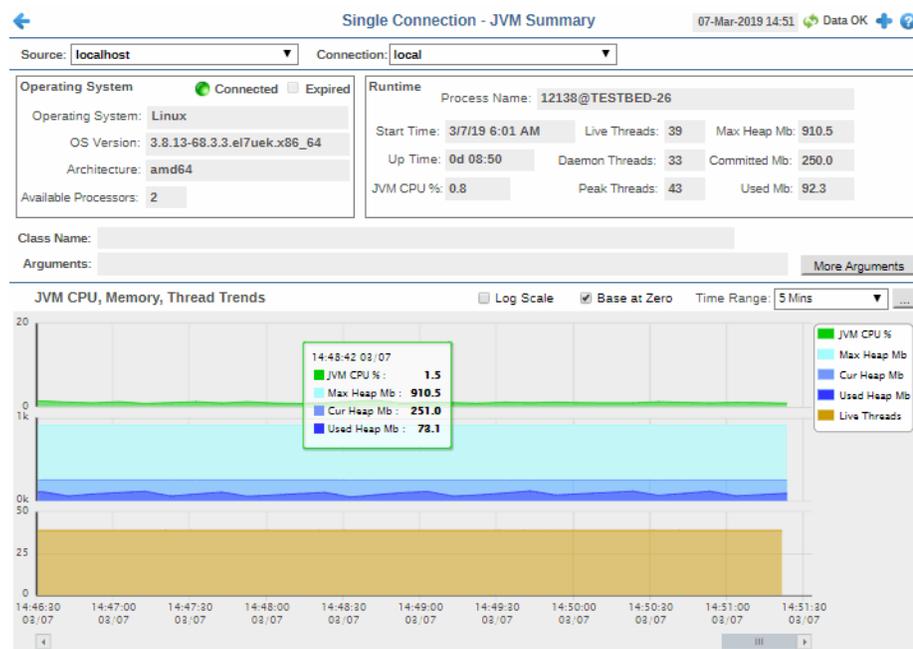
Use these detailed JVM displays to investigate performance issues on a JVM. Displays in this View are:

- **"JVM Summary"**: Table of connection details for a single JVM as well as performance trend graphs.
- **"JVM System Properties"**: Table of system details for a single JVM.
- **"JVM Memory Pool Trends"**: Trend graphs of memory pool utilization.
- **"JVM GC Trends"**: Trend graphs of garbage collection memory utilization.

## JVM Summary

Track JVM memory and CPU usage, get JVM system information, application performance metrics, and input arguments for a single connection. Verify whether the memory usage has reached a plateau. Or, if usage is getting close to the limit, determine whether to allocate more memory.

Use the available drop-down menus or right-click to filter data shown in the display.



### Title Bar (possible features are):

- ← ↑ Open the previous and upper display.
- + Open an instance of this display in a new window.
- ? Open the online help page for this display.
- Menu Table open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

**Fields and Data**

This display includes:

- Source** Select the type of connection to the RTView Server.
- Connection** Select an RTView Server from the drop-down menu. Names can be modified in the RTView Server configuration properties file.
- Operating System**  
Displays data pertaining to the operating system running on the host on which the JVM resides.
- |                             |   |
|-----------------------------|---|
| <b>Connected</b>            | The data connection state:<br> Disconnected.<br> Connected. |
| <b>Expired</b>              | When checked, this server is expired due to inactivity.   |
| <b>Operating System</b>     | The name of the operating system running on the host on which the JVM resides.  |
| <b>OS Version</b>           | The operating system version.   |
| <b>Architecture</b>         | The ISA used by the processor.  |
| <b>Available Processors</b> | The total number of processors available to the JVM.  |

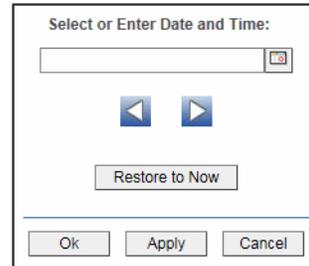
**Runtime**

<b>Process Name</b>	Name of the process.
<b>Start Time</b>	The date and time that the application started running.
<b>Up Time</b>	The amount of time the application has been running, in the following format: <b>0d 00:00</b> <b>&lt;days&gt;d &lt;hours&gt;:&lt;minutes&gt;:&lt;seconds&gt;</b> For example: <b>10d 08:41:38</b>
<b>JVM CPU %</b>	The amount of CPU usage by the JVM, in percent.
<b>Live Threads</b>	The total number of live threads.
<b>Daemon Threads</b>	The total number of live daemon threads.
<b>Peak Threads</b>	The total number of peak live threads since the JVM started or the peak was reset.
<b>Max Heap Mb</b>	The maximum amount of memory used for memory management by the application in the time range specified. This value may change or be undefined.  NOTE: A memory allocation can fail if the JVM attempts to set the <b>Used</b> memory allocation to a value greater than the <b>Committed</b> memory allocation, even if the amount for <b>Used</b> memory is less than or equal to the <i>Maximum</i> memory allocation (for example, when the system is low on virtual memory).
<b>Committed Mb</b>	The amount of memory, in megabytes, guaranteed to be available for use by the JVM. The amount of committed memory can be a fixed or variable size. If set to be a variable size, the amount of committed memory can change over time, as the JVM may release memory to the system. This means that the amount allocated for <b>Committed</b> memory could be less than the amount initially allocated. <b>Committed</b> memory will always be greater than or equal to the amount allocated for <b>Used</b> memory.
<b>Used Mb</b>	The amount of memory currently used by the application. Memory used includes the memory occupied by all objects including both reachable and unreachable objects.
<b>Class Name</b>	Class name used for JVM.
<b>Arguments</b>	The arguments used to start the application.
<b>More Arguments</b>	Additional arguments used to start the application.

**JVM CPU, Memory, Thread Trends**

Shows JVM metrics for the selected server.

- Log Scale** Enable to use a logarithmic scale for the Y axis. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Base at Zero** Use zero as the Y axis minimum for all graph traces.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

- JVM CPU %** Traces the amount of memory, in percent, used by the JVM in the time range specified.
- Max Heap Mb** Traces the maximum amount of memory used for memory management by the application in the time range specified. This value may change or be undefined.  
NOTE: A memory allocation can fail if the JVM attempts to set the **Used** memory allocation to a value greater than the **Committed** memory allocation, even if the amount for **Used** memory is less than or equal to the **Maximum** memory allocation (for example, when the system is low on virtual memory).
- Cur Heap Mb** Traces the current amount of memory, in megabytes, used for memory management by the application in the time range specified.
- Used Heap Mb** Traces the memory currently used by the application.
- Live Threads** Traces the total number of currently active threads in the time range specified.

## JVM System Properties

Track JVM input arguments and system properties for a single connection. Use the available drop-down menus or right-click to filter data shown in the display.

Property	Value
awt.toolkit	sun.awt.X11.XToolkit
com.sl.rtvie.customRtvAppManagerClassName	com.sl.gmsjrtvutils.RtvAppManager
com.sl.rtvie.log4jFile	/home/m/testbed/TBSOLMON/RTView/SolaceMonitor/rtvapm/common/conf/sl.log
com.sl.rtvie.RTVLog4jLevel	info
com.sl.rtvie.showLogCategory	true
com.sl.rtvie.useLog4j	true
com.sun.management.jmxremote.authenticate	false
com.sun.management.jmxremote.port	3068
com.sun.management.jmxremote.ssl	false
file.encoding	UTF-8
file.encoding.pkg	sun.io
file.separator	/
java.awt.graphicsenv	sun.awt.X11GraphicsEnvironment
java.awt.printerjob	sun.print.PSPrinterJob

### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu , Table open commonly accessed displays.

Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green Data OK icon is a strong indication that data is current and valid.

### Fields and Data

This display includes:

- Source** Select the type of connection to the RTView Server.
- Connection** Select an RTView Server from the drop-down menu. Names can be modified in the RTView Server configuration properties file.
- Connected** The data connection state:
  - Disconnected.
  - Connected.
- Java Version** The Java version running on the selected server.
- JVM Arguments** The JVM arguments in the **RuntimeMXBean InputArguments** attribute.
- Command Line Arguments** Arguments used to start the application.

### System Properties

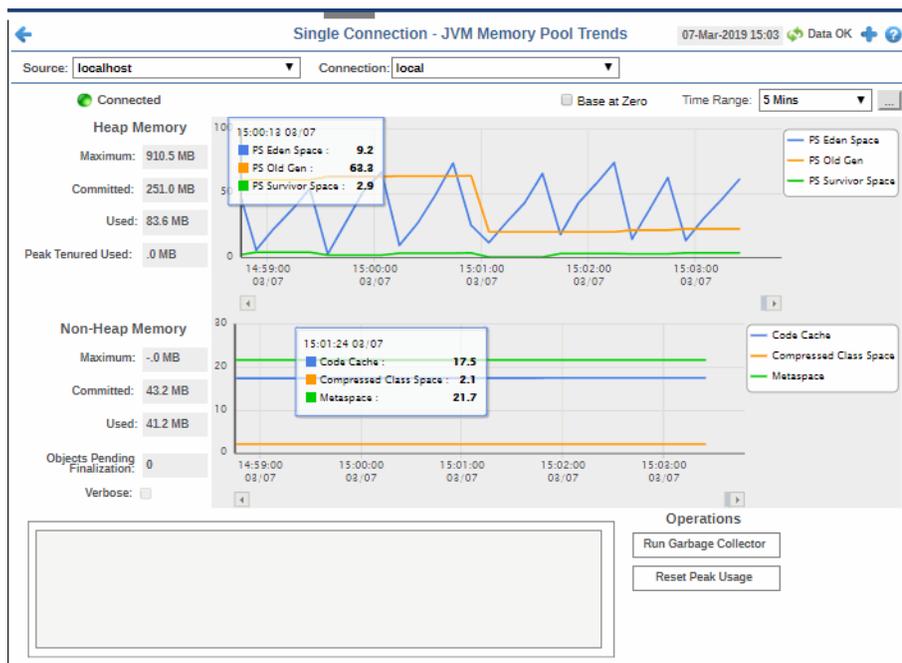
This table lists and describes system property settings.

**Property** Name of the property.

**Value** Current value of the property.

### JVM Memory Pool Trends

Track JVM heap and non-heap memory usage for a single connection. Use the available drop-down menus or right-click to filter data shown in the display.



#### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu** , **Table** open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

### Fields and Data

This display includes:

**Source** Select the type of connection to the RTView Server.

**Connection** Select an RTView Server from the drop-down menu. Names can be modified in the RTView Server configuration properties file.

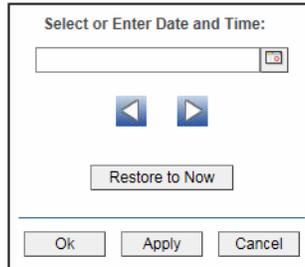
**Connected** The data connection state:

● Disconnected.

● Connected.

**Base at Zero** Use zero as the Y axis minimum for all graph traces.

**Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

## Heap Memory

<b>Maximum</b>	<p>The maximum amount of memory used, in megabytes, for memory management by the application in the time range specified. This value may change or be undefined.</p> <p>NOTE: A memory allocation can fail if the JVM attempts to set the <b>Used</b> memory allocation to a value greater than the <b>Committed</b> memory allocation, even if the amount for <b>Used</b> memory is less than or equal to the <b>Maximum</b> memory allocation (for example, when the system is low on virtual memory).</p>
<b>Committed</b>	<p>The amount of memory, in megabytes, guaranteed to be available for use by the JVM. The amount of committed memory can be a fixed or variable size. If set to be a variable size, the amount of committed memory can change over time, as the JVM may release memory to the system. This means that the amount allocated for <b>Committed</b> memory could be less than the amount initially allocated. <b>Committed</b> memory will always be greater than or equal to the amount allocated for <b>Used</b> memory.</p>
<b>Used</b>	<p>The amount of memory, in megabytes, currently used by the application. Memory used includes the memory occupied by all objects including both reachable and unreachable objects.</p>
<b>Peak Tenured Used</b>	<p>The amount of memory, in megabytes, used by tenured JVM objects in the time range specified. Tenured refers to JVM objects contained in a pool that holds objects that have avoided garbage collection and reside in the survivor space. Peak tenured refers to the maximum value of the tenured memory over a specified period of time.</p>
<b>Eden Space</b>	<p>Traces the amount of memory used by the JVM eden pool in the time range specified. Eden refers to the JVM eden pool, which is used to initially allocate memory for most objects.</p>
<b>Survivor Space</b>	<p>Traces the amount of memory used by the JVM survivor pool in the time range specified. The JVM survivor pool holds objects that survive the eden space garbage collection.</p>
<b>Tenured Gen</b>	<p>Traces the amount of memory used by tenured JVM objects in the time range specified. Tenured refers to JVM objects contained in a pool that holds objects that have avoided garbage collection and reside in the survivor space. Peak tenured refers to the maximum value of the tenured memory over a specified period of time.</p>

### Non-Heap Memory

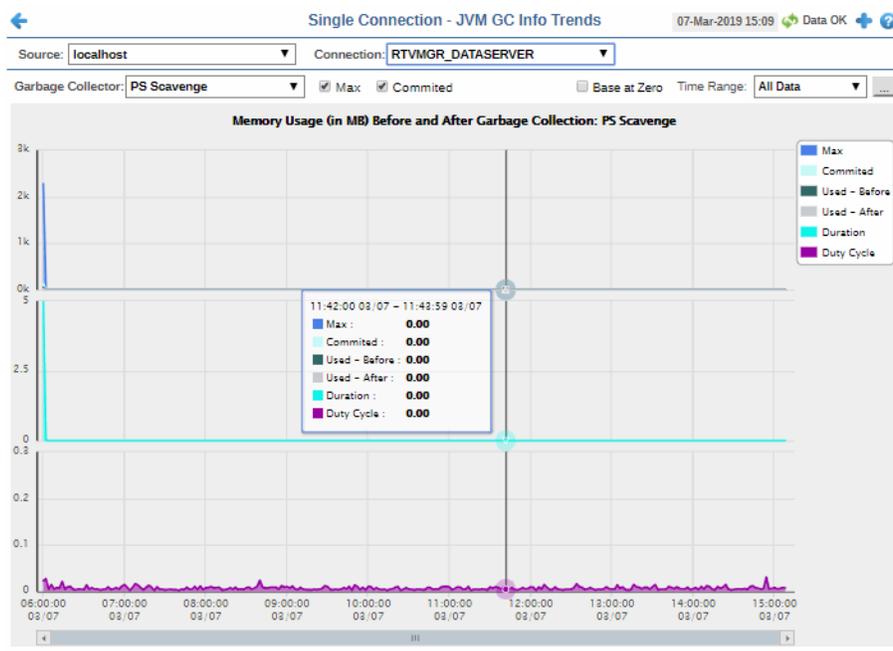
<b>Maximum</b>	The maximum amount of memory, in megabytes, used for JVM non-heap memory management by the application in the time range specified.
<b>Committed</b>	The amount of memory, in megabytes, guaranteed to be available for use by JVM non-heap memory management. The amount of committed memory can be a fixed or variable size. If set to be a variable size, it can change over time, as the JVM may release memory to the system. This means that the amount allocated for <b>Committed</b> memory could <b>be</b> less than the amount initially allocated. Committed memory will always be greater than or equal to the amount allocated for <b>Used</b> memory.
<b>Used</b>	The amount of memory, in megabytes, currently used by the application. Memory used includes the memory occupied by all objects including both reachable and unreachable objects.
<b>Objects Pending Finalization</b>	The value of the <b>MemoryMXBean ObjectPendingFinalizationCount</b> attribute.
<b>Verbose</b>	The value of the <b>MemoryMXBean Verbose</b> attribute.
<b>Code Cache</b>	Traces the amount of non-heap memory used in the JVM for compilation and storage of native code.
<b>Perm Gen</b>	Traces the amount of memory used by the pool containing reflective data of the virtual machine, such as class and method objects. With JVMs that use class data sharing, this generation is divided into read-only and read-write areas.

### Operations

<b>Run Garbage Collector</b>	Performs garbage collection on the selected server.
<b>Reset Peak Usage</b>	Clears peak usage on the selected server.

## JVM GC Trends

Track JVM garbage collection memory usage for a single connection. Use the available drop-down menus or right-click to filter data shown in the display.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

### Fields and Data

This display includes:

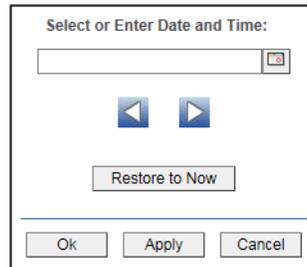
- Source** Select the type of connection to the RTView Server.
- Connection** Select an RTView Server from the drop-down menu. Names can be modified in the RTView Server configuration properties file.
- Garbage Collector** Select a garbage collection method: **Copy** or **MarkSweepCompact**.
- Max** Shows the maximum amount of memory used for JVM garbage collection in the time range specified.
- Committed** Shows the amount of memory guaranteed to be available for use by JVM non-heap memory management. The amount of committed memory can be a fixed or variable size. If set to be a variable size, it can change over time, as the JVM may release memory to the system. This means that the amount allocated for **Committed** memory could be less than the amount initially allocated. **Committed** memory will always be greater than or equal to the amount allocated for **Used** memory.

**Base at Zero**

Use zero as the Y axis minimum for all graph traces.

**Time Range**

Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .



By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

**Memory Usage (in MB) Before and After Garbage Collection**

<b>Maximum</b>	Traces the maximum amount of memory used by garbage collection in the time range specified. This value may change or be undefined.  NOTE: A memory allocation can fail if the JVM attempts to set the <b>Used</b> memory allocation to a value greater than the <b>Committed</b> memory allocation, even if the amount for <b>Used</b> memory is less than or equal to the <b>Maximum</b> memory allocation (for example, when the system is low on virtual memory).
<b>Committed</b>	Traces the amount of memory guaranteed to be available for use by the JVM. The amount of committed memory can be a fixed or variable size. If set to be a variable size, the amount of committed memory can change over time, as the JVM may release memory to the system. This means that the amount allocated for <b>Committed</b> memory could be less than the amount initially allocated. <b>Committed</b> memory will always be greater than or equal to the amount allocated for <b>Used</b> memory.
<b>Used - Before</b>	Traces the amount of memory used before the last garbage collection.
<b>Used - After</b>	Traces the amount of memory used after the last garbage collection.
<b>Duration</b>	The duration, in seconds, of garbage collection.
<b>Duty Cycle</b>	The percentage of time that the application spends in garbage collection.

---

## JVM Displays - HTML

The RTView Manager JVM displays present performance data for monitored Java Virtual Machine (JVM) processes. Use these displays to examine the current and historical performance metrics and resource usage of JVMs. Any JVM that is enabled for monitoring can be included in these displays. The displays include summary overviews and detail pages with historical trends.

You can set alert thresholds on performance and resource metrics for your JVMs, including **CPU Percent**, **Memory Used** and **Gc Duty cycle**. Alerts are shown in the ["JVMs Heatmap - HTML"](#) display. Use the detailed JVM displays to investigate further; for example a **Memory Used** alarm might take you to the ["JVM Summary - HTML"](#) display to get historical memory use, or a **Gc Duty Cycle** alarm might take you to the ["JVM GC Trends - HTML"](#) display. Displays in this View are:

The HTML version features an overview display, ["JVM Overview - HTML"](#) (pictured below), and the following displays which can be found under **Components** tab > **Processes / JVM Processes** once RTView Manager is installed:

- ["JVMs Heatmap - HTML"](#): Heatmap of alert states for all JVM connections
- ["JVM Summary - HTML"](#): Table of connection details for all JVM connections.
- ["JVM System Properties - HTML"](#): Table of connection details for a single JVM as well as performance trend graphs.
- ["JVM GC Trends - HTML"](#): Table of system details for a single JVM.
- ["JVMs Heatmap - HTML"](#): Trend graphs of memory pool utilization.
- ["JVM GC Trends - HTML"](#): Trend graphs of garbage collection memory utilization.

### JVM Overview - HTML

The **JVM Overview** is the top-level display for the JVM Solution Package, which provides a good starting point for immediately getting the status of all your JVM instances on your Data Server.

Choose a DataServer for which you want to see data and easily view the current data for that DataServer including:

- The total number of active alerts, including the total number of critical and warning alerts.
- The number of JVMs and the **Top CPU %** user across all servers.
- The maximum memory used and maximum number of threads.

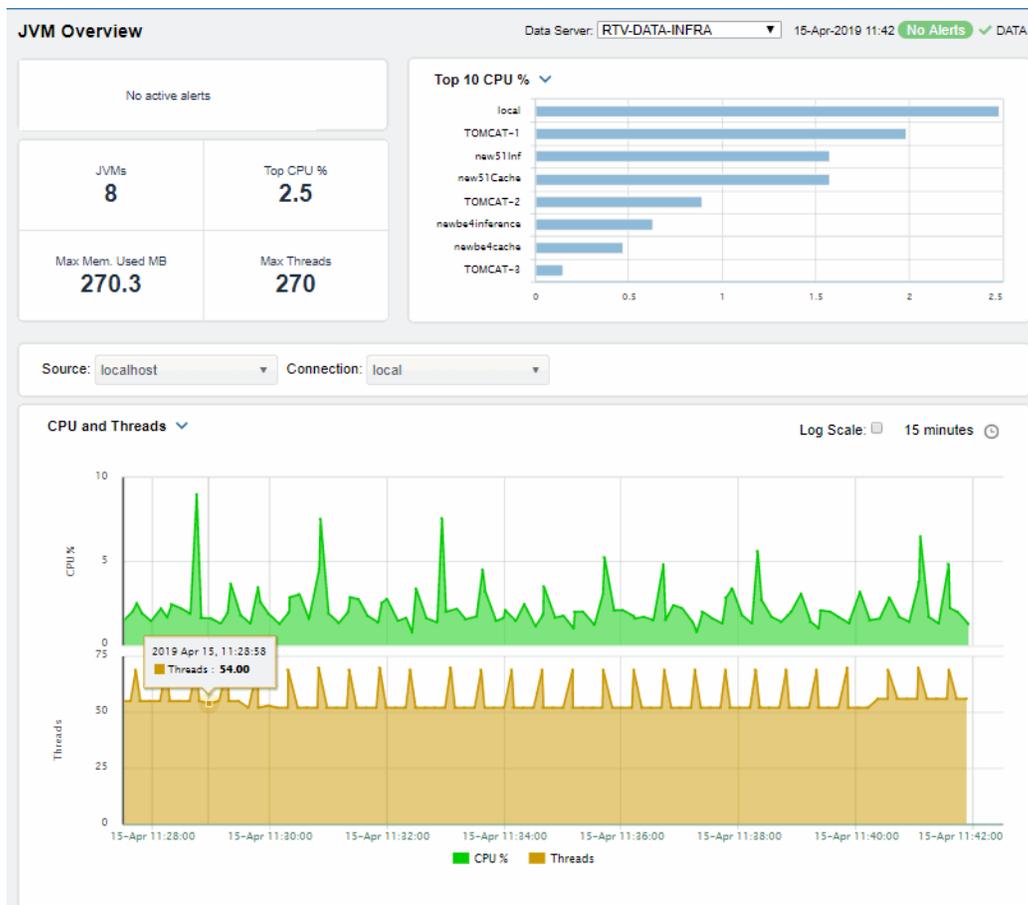
A bar graph shows the JVMs with **Top 10 CPU %** utilization. Use the drop-down menu to show JVMs with **Top 10 Used Heap** memory utilization and JVMs with **Top 10 Live Threads**.

You can hover over each region in the upper half of the Overview to see more detail. You can also drill down to see even more detail by clicking on each respective region in the Overview.

For example, clicking on the alerts in the CRITICAL and WARNING alerts region opens the Alerts Table display. Clicking on **Top CPU %** opens the ["JVM Summary - HTML"](#) display.

The bottom half of the display provides a performance trend graph for a connection on the DataServer. Choose a **Source** and **Connection** from the drop-down menus. Use the trend graph drop-down menu to show metrics for **CPU and Threads** utilization or Heap Memory utilization.

You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.



## JVMs Table - HTML

Investigate JVM connection utilization metrics and configuration information for one or all JVMs. Choose one or **All** JVMs from the **Source** drop-down menu. Each row in the table contains data for a particular connection on the selected JVM(s).

This display contains all metrics available for JVM connections, including the **Port** number and the current most critical **Alert Level**, where:

- Red indicates that one or more alerts exceeded their ALARM LEVEL threshold in the table row.
- Yellow indicates that one or more alerts exceeded their WARNING LEVEL threshold in the table row.
- Green indicates that no alerts exceeded their WARNING or ALARM LEVEL threshold in the table row.

Double-click on a table row to drill-down to the **JVM Summary - HTML** display and view metrics for the JVM hosting the connection.

Check the **Show Inactive** box to include inactive connections.

Click a column header to sort column data in ascending or descending order. Toggle between the commonly accessed **Table** and **Heatmap** displays by clicking the drop down list on the display title.

**JVMs Table** 15-Apr-2019 12:59 No Alerts  DATA

Source: - All -

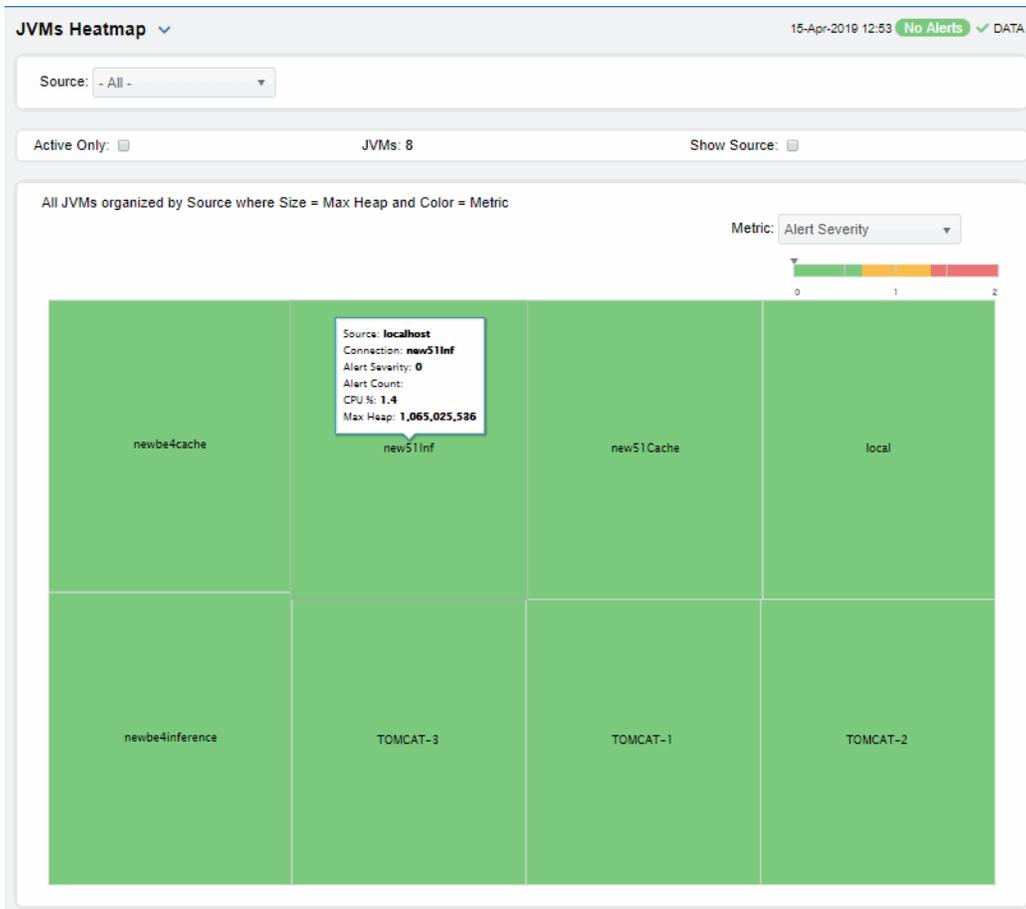
Active Only:  JVMs: 14

Source	Connection	Expired	Connected	Alert Level	Alert Count	Host	Port	PID	CPU %
localhost	TOMCAT-3		✓	✓		172.30.1.85	9999		0.
localhost	TOMCAT-2		✓	✓		172.30.1.57	9999		1.
localhost	TOMCAT-1		✓	✓		172.30.1.174	9999		1.
localhost	local		✓	✓				348@172.30.1.104	1.
localhost	local		✓	✓				9@172.30.1.174	1.
localhost	local		✓	✓				1114@172.30.1.31	1.
localhost	local		✓	✓				1120@172.30.1.31	1.
localhost	local		✓	✓				123@172.30.1.97	1.
localhost	local		✓	✓				126@172.30.1.97	1.
localhost	local		✓	✓				9@172.30.1.174	1.
localhost	new51Inf		✓	✓		192.168.200.144	58701		1.
localhost	newbe4cache		✓	✓		192.168.200.144	48700		0.
localhost	new51Cache		✓	✓		192.168.200.144	58700		0.
localhost	newbe4inference		✓	✓		192.168.200.144	48701		0.

## JVMs Heatmap - HTML

View the most critical alert state for all monitored JVM connections for one or all sources, as well as CPU and memory utilization. The heatmap organizes JVM connections by source and host, and uses color to show the most critical Metric value for each JVM connection associated with the selected source. Each rectangle in the heatmap represents a JVM connection. The rectangle size represents the amount of memory reserved for that process; a larger size is a larger value. Each Metric (selected from the drop-down menu) has a color gradient bar that maps relative values to colors.

Choose one or **All Sources** from the **Source** drop-down menu. Use the check-boxes  to include or exclude labels in the heatmap. Move your mouse over a rectangle to see detailed JVM connection information (including **PID**). Drill-down and investigate by clicking a rectangle in the heatmap to view details for the selected connection in the **JVM Summary** display.



### Metric

Select the Metric to display in the heatmap. Each Metric has a color gradient bar that maps relative values to colors.

#### Alert Severity

The maximum level of alerts in the heatmap rectangle. Values range from **0** - **2**, as indicated in the color gradient bar, where **2** is the highest Alert Severity.

Red indicates that one or more alerts have reached their alarm threshold. Alerts that have exceeded their specified ALARM LEVEL threshold have an Alert Severity value of **2**.

Yellow indicates that one or more alerts have reached their alarm threshold. Alerts that have exceeded their specified WARNING LEVEL threshold have an Alert Severity value of **1**.

Green indicates that no alerts have reached their alert thresholds. Alerts that have not exceeded their specified thresholds have an Alert Severity value of **0**.

#### Alert Count

The number of alerts for the rectangle. The color gradient bar values range from **0** to the maximum number of alerts in the heatmap.

<b>CPU %</b>	The total percent (%) CPU utilization for the rectangle. The color gradient  bar values range from <b>0</b> to the maximum percent (%) CPU utilization in the heatmap.
<b>Memory %</b>	The total percent (%) memory utilization for the rectangle. The color gradient  bar values range from <b>0</b> to the maximum percent (%) memory utilization in the heatmap.
<b>Current Heap</b>	The current amount of heap committed for the connection, in kilobytes. The color gradient  bar values range from <b>0</b> to the maximum amount in the heatmap.
<b>Used Heap</b>	The total amount of heap used by the connection, in kilobytes. The color gradient  bar values range from <b>0</b> to the maximum amount used in the heatmap.

## JVM Summary - HTML

Track utilization by a single connection on a JVM, including **Memory** and **CPU** usage, amount of **Committed Memory** (the amount of memory, in megabytes, guaranteed to be available for use by the JVM. The amount of committed memory can be a fixed or variable size. If set to be a variable size, the amount of committed memory can change over time, as the JVM may release memory to the system. This means that the amount allocated for **Committed** memory could be less than the amount initially allocated. Committed memory will always be greater than or equal to the amount allocated for **Used memory**) and **Maximum Memory** used, number of **Threads** and **Peak Threads**. Verify whether the memory usage has reached a plateau. Or, if usage is getting close to the limit, determine whether to allocate more memory.

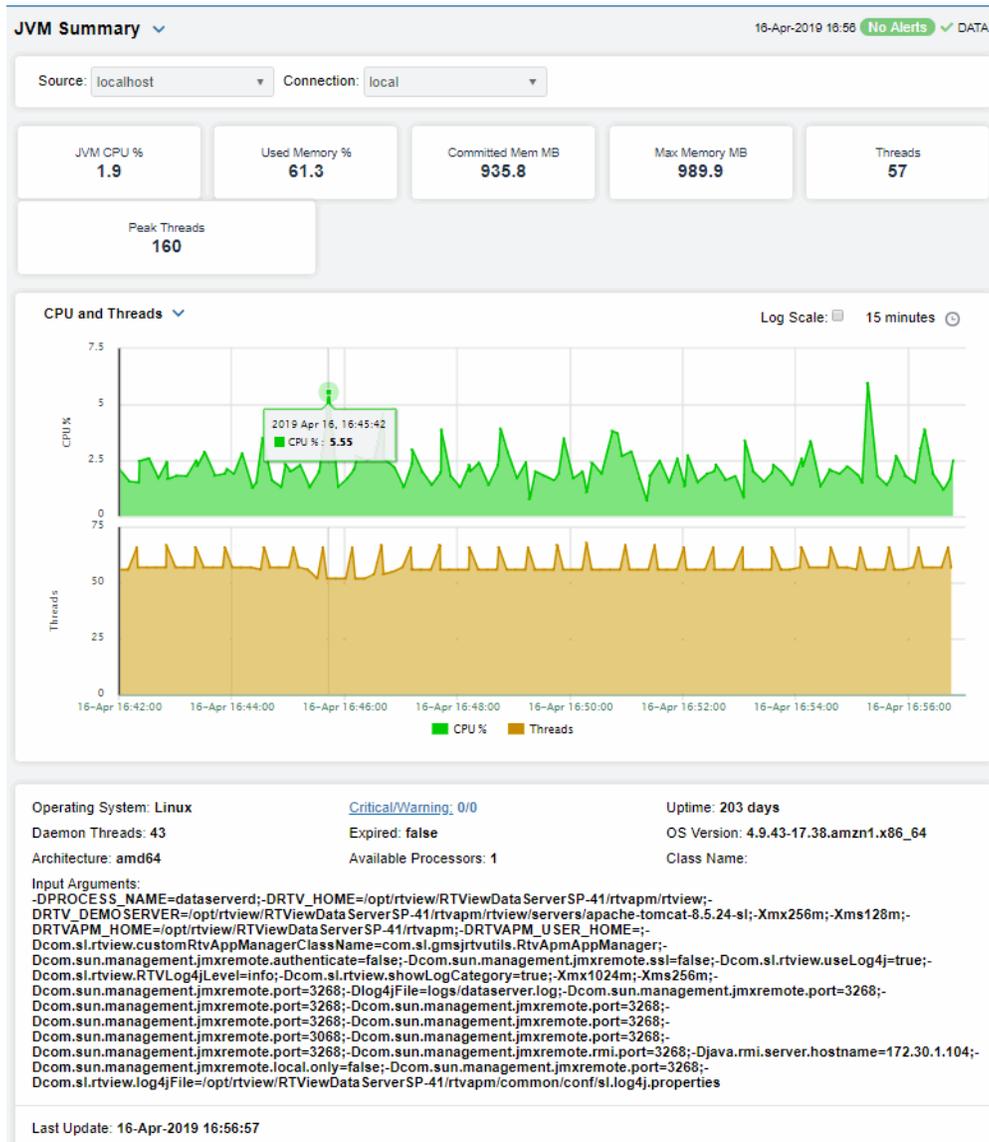
Clicking on the information boxes at the top of the display takes you to the **JVMs Table - HTML** display, where you can view and compare with other connections on the JVM.

You can set the time range for the trend graph to trace. You can also choose what to trace from the drop-down menu:

- **CPU and Threads** traces the amount of CPU used by the JVM and the total number of live threads.
- **Heap Memory** traces the amount of memory used for memory management by the application in the time range specified. This value may change or be undefined. Note that a memory allocation can fail if the JVM attempts to set the Used memory allocation to a value greater than the **Committed** memory allocation, even if the amount for **Used** memory is less than or equal to the Maximum memory allocation (for example, when the system is low on virtual memory).

At the bottom of the display you also can get JVM operating system information, the number of processors available to the JVM, the **Architecture** which is the ISA used by the processor, the number of **Daemon Threads** and **Input Arguments** for starting JVM.

Clicking the **Critical/Warning** link at the bottom of the display opens the Alerts Table by Component display.



## JVM System Properties - HTML

View JVM arguments in the RuntimeMXBean InputArguments attribute, command line arguments for starting applications and system properties settings for a connection. Choose a **Source** and **Connection** from the drop-down menus.

Click a column header to sort column data in ascending or descending order or right-click to filter data shown in the display. Toggle between the commonly accessed Table and Heatmap displays by clicking the drop down list on the display title.

**JVM System Properties** 15-Apr-2019 13:05 No Alerts DATA

Source: localhost Connection: local

**JVM Arguments**

Value
-Dcom.slr.view.customRtvAppManagerClassName=com.slr.gmsjrtutils.RtvAppManager
-Dcom.sun.management.jmxremote.authenticate=false
-Dcom.sun.management.jmxremote.ssl=false
-Dcom.slr.view.useLog4j=true
-Dcom.slr.view.RTVLog4jLevel=info
-Dcom.slr.view.showLogCategory=true
-Xmx1024m
-Xms256m

**Command-Line Arguments**

Value
-------

**System Properties**

Property	Value
com.sun.management.jmxremote.rmi.port	3288
awt.toolkit	sun.awt.windows.WToolkit
com.slr.view.RTVLog4jLevel	info
file.encoding.pkg	sun.io
java.specification.version	1.7
sun.cpu.isalist	amd64

## JVM GC Trends - HTML

Track JVM garbage collection memory utilization trends for a single connection. Choose a **Source**, **Connection** and **Garbage Collector** (options are **Copy** or **MarkSweepCompact**) from the drop-down menus. The upper trend graph traces the following for the selected garbage collector for the time range specified:

- **Max:** The maximum amount of memory, in megabytes, used for JVM garbage collection.
- **Committed:** The amount of memory, in megabytes, guaranteed to be available for use by JVM non-heap memory management. Note that the amount of committed memory can be a fixed or variable size. If set to be a variable size, it can change over time, as the JVM may release memory to the system. This means that the amount allocated for committed memory could be less than the amount initially allocated. Committed memory will always be greater than or equal to the amount allocated for used memory.
- **Used - Before:** The amount of memory, in megabytes, used before the last garbage collection.
- **Used - After:** The amount of memory, in megabytes, used after the last garbage collection.

The lower trend graph traces the following for the selected garbage collector for the time range specified:

- **Duration:** The duration, in seconds, of garbage collection.
- **Duty Cycle:** The percentage of time that the application spends in garbage collection.

You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.



---

## Tomcat Displays

The Tomcat displays provide extensive visibility into the health and performance of Tomcat application servers and installed web modules. The following Tomcat Views (and their associated displays) can be found under **Components** tab > **Application/Web Servers** > **Tomcat**. The Tomcat displays come with RTView Enterprise.

Tomcat has the following Views:

- [“Tomcat Servers”](#)
- [“Tomcat Applications”](#)

### Tomcat Servers

These displays present performance data for monitored Tomcat Servers. Use these displays to examine the state and performance of your Tomcat servers as well as all installed web modules. The server displays include summary overviews and detail pages with historical trends. Displays in this View are:

- [“All Tomcat Servers”](#): Table of connection details and performance metrics for all Tomcat connections.
- [“Tomcat Server Summary”](#): Performance metrics for one Tomcat Server, including current and historic performance metrics.

#### All Tomcat Servers

View Tomcat Server details per connection such as the total number of sessions, bytes sent/received, and processing time. Each row in the table is a different Tomcat Server. The row color for inactive connections is dark red.

Use this display to get Tomcat server session counts, access and request rates, cache hit rates and data transmission metrics.

Drill-down and investigate by clicking a row in the table to view details for the selected connection in the **Service Summary** display.

Connection	Source	Sessions Active	Sessions Total	Sessions Expired	Accesses per sec	Accesses Total	Bytes Rcvd per sec	Bytes Rcvd Total
TOMCAT	localhost	4	17	13	1.4	30,302	603.1	433,851.8

**Title Bar** (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

**Fields and Data**

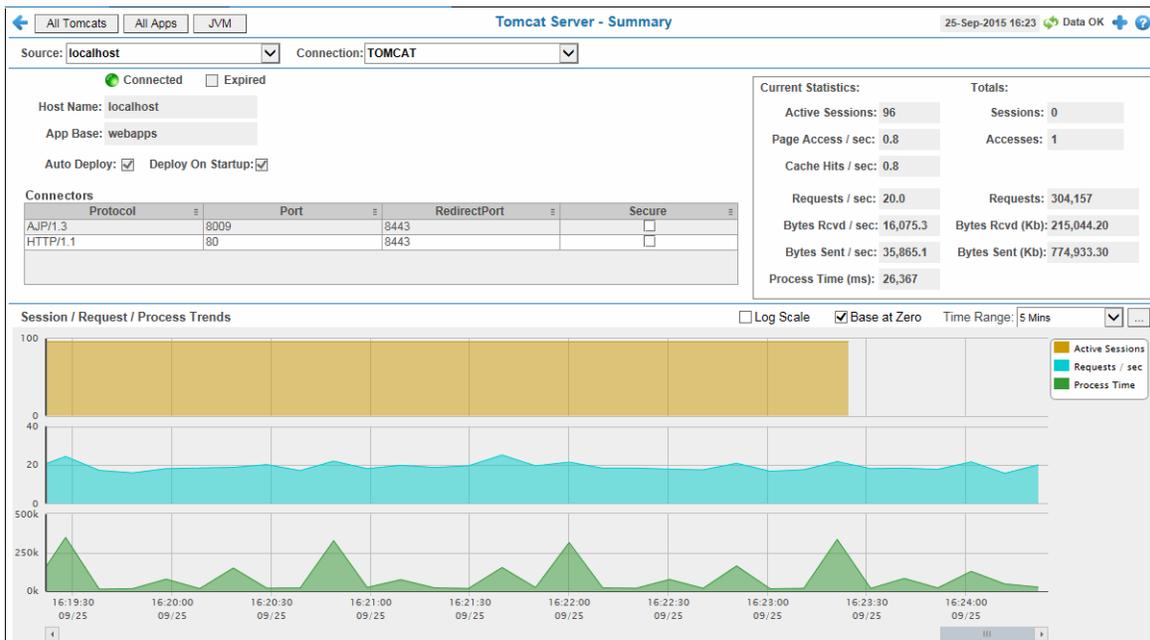
This display includes:

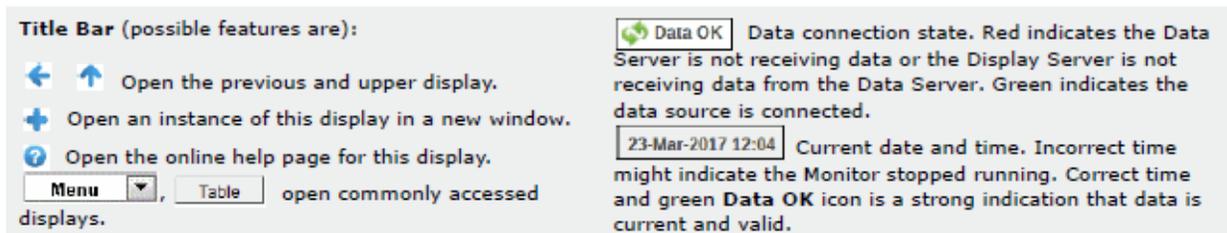
<b>Tomcat Count</b>	The number of Tomcat connections in the table.
<b>Connection</b>	The name of the Tomcat connection.
<b>Source</b>	The host where the Tomcat Server is running.
<b>Sessions Active</b>	The number of currently active client sessions.
<b>Sessions Total</b>	The total number of client sessions since the server was started.
<b>Sessions Expired</b>	The total number of client sessions that expired since the server was started.
<b>Accesses per sec</b>	The number of times pages are accessed, per second.
<b>Accesses Total</b>	The total number of times pages have been accessed since the server was started.
<b>Bytes Rcvd per sec</b>	The number of bytes received per second.
<b>Bytes Rcvd Total</b>	The total number of bytes received since the server was started.

- Bytes Sent per sec**      The number of bytes sent per second.
- Bytes Sent Total**      The total number of bytes sent since the server was started.
- Cache Hit Rate**      The number of times the cache is accessed, per second.
- Requests per sec**      The number of requests received, per second.
- Requests Total**      The total number of requests received since the server was started.
- Process Time**      The average amount of time, in milliseconds, to process requests.
- Error Count**      The number of errors that have occurred since the server was started.
- appBase**      The directory in which Tomcat is installed.
- Display Name**      The name of the currently open display.
- Expired**      When checked, this connection is expired due to inactivity.
- time\_stamp**      The date and time this row of data was last updated.  
Format:  
**MM/DD/YY HH:MM:SS**  
**<month>/ <day>/<year> <hours>:<minutes>:<seconds>**

### Tomcat Server Summary

Track the performance of one Tomcat Server and get Tomcat hosting and connection details. You can drill down to this display from the Servers table for detailed information and historical trends for a specific server. The trends include Active Sessions, Requests per Sec, and Process Time.





### Fields and Data

This display includes:

- Source** Select the host where the Tomcat Server is running.
- Connection** Select a Tomcat Server from the drop-down menu.
- Connected** The Tomcat Server connection state:
  - Disconnected.
  - Connected.
- Expired** When checked, this server is expired due to inactivity.
- Host Name** The name of the host where the application resides.
- App Base** The directory in which Tomcat modules are installed.
- Auto Deploy** When checked, indicates that the Tomcat option, automatic application deployment, is enabled.
 

**Note:** This Tomcat option is set using the **autoDeploy** property in the **server.xml** file, located in the Tomcat **conf** directory. **autoDeploy=true** enables the option.
- Deploy On Startup** When checked, indicates that the option to deploy the application on Tomcat startup is enabled.
 

**Note:** This Tomcat option is set using the **deployOnStartup** property in the **server.xml** file, located in the Tomcat **conf** directory. When enabled (**deployOnStartup=true**), applications from the host are automatically deployed.

### Connectors

This table shows Tomcat application connection information.

<b>Protocol</b>	The protocol used by the Tomcat application on the host.
<b>Port</b>	The port number used by the Tomcat application on the host.
<b>RedirectPort</b>	The redirect port number used by the Tomcat application on the host.
<b>Secure</b>	When checked, specifies that the Tomcat application uses a secure connection on the host.

### Current Statistics / Totals

<b>Active Sessions</b>	The number of clients currently in session with the servlet.
<b>Sessions</b>	The total number of client sessions since the server was started.
<b>Page Access / sec</b>	The number of times pages are accessed, per second.
<b>Accesses</b>	The total number of page accesses since the server was started.
<b>Cache Hits / sec</b>	The number of times the cache is accessed, per second.
<b>Requests / sec</b>	The number of requests received, per second.
<b>Requests</b>	The total number of requests since the server was started.
<b>Bytes Rcvd / sec</b>	The number of bytes received, per second.
<b>Bytes Rcvd (Kb)</b>	The number of kilobytes received since the server was started.
<b>Bytes Sent / sec</b>	The number of bytes sent, per second.
<b>Bytes Sent (Kb)</b>	The total number of kilobytes sent since the server was started.
<b>Process Time</b>	The amount of time, in milliseconds, for the servlet to process client requests.

**Session / Request / Process Trends**

Shows metrics for the selected server.

- Log Scale** Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Base at Zero** Use zero as the Y axis minimum for all graph traces.
- Time Range** Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .

By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

- Active Sessions** Traces the number of currently active client sessions.
- Requests /sec** Traces the number of requests received, per second.
- Process Time** Traces the average amount of time, in milliseconds, to process requests.

## Tomcat Applications

These displays present performance data for monitored Tomcat Applications. Use these displays to examine the state and performance of your Tomcat applications and all installed web modules. The server displays include summary overviews and detail pages with historical trends.

Displays in this View are:

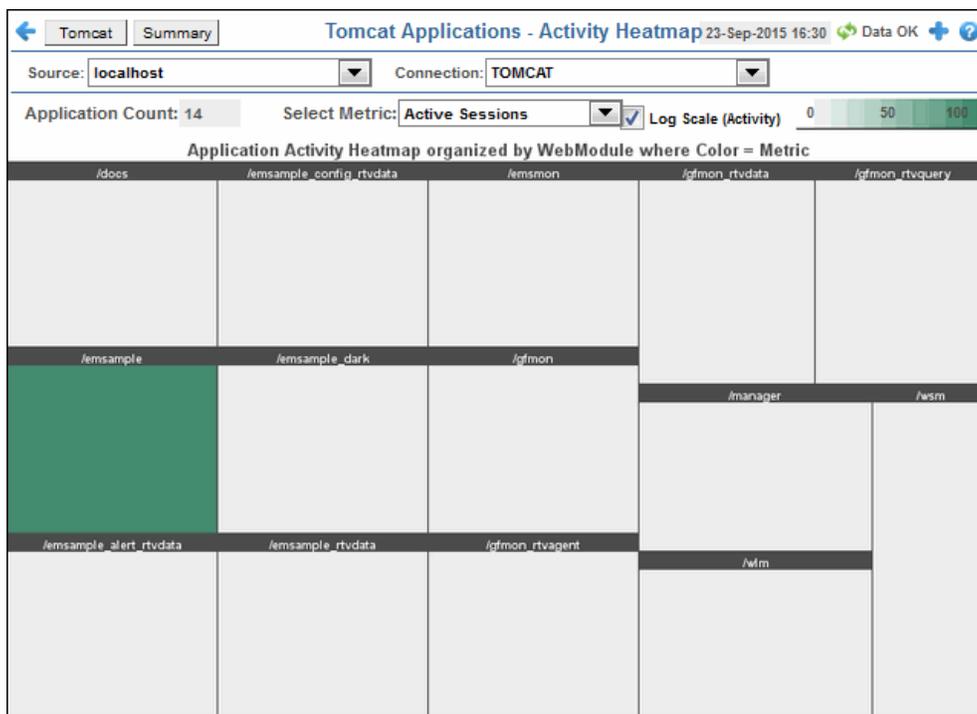
- **"Applications Heatmap"**: Heatmap of performance metrics for all Web modules for one Tomcat Server.
- **"Applications Summary"**: Table and trend graphs of performance metrics for Web modules.

## Applications Heatmap

View performance metrics for all monitored Tomcat Web modules for one Tomcat Server. The heatmap organizes Tomcat Web modules by server, and uses color to show the most critical Metric value for each Tomcat connection associated with the selected source. Each rectangle in the heatmap represents a Web module. In this heatmap, the rectangle size is the same for all Web modules. Each Metric (selected from the drop-down menu) has a color gradient bar that maps relative values to colors.

Use this display to see at-a-glance the health of all your web applications. You can select the heatmap color metric from a list including active sessions, access rate, and total access count.

Use the available drop-down menus or right-click to filter data shown in the display. Use the check-boxes  to include or exclude labels in the heatmap. Move your mouse over a rectangle to see additional information. Drill-down and investigate by clicking a rectangle in the heatmap to view details for the selected Web module in the **Application Summary** display.



### Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- open commonly accessed displays.

**Data OK** Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.

**23-Mar-2017 12:04** Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.

### Fields and Data

This display includes:

- Source** Select the host where the Tomcat Server is running.
- Connection** Select a Tomcat Server from the drop-down menu.

**Application Count** The number of Tomcat applications in the heatmap.

**Log Scale (Activity)** Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.

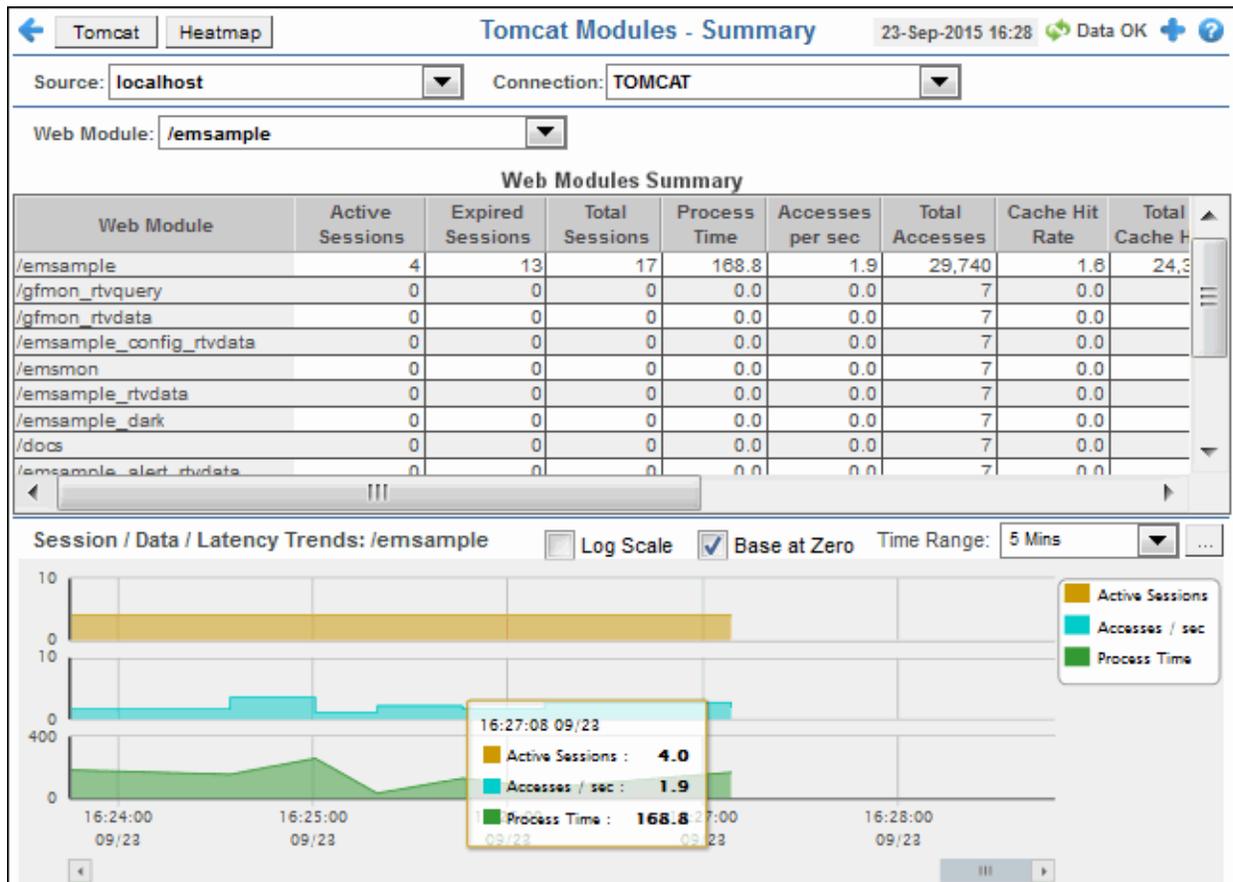
**Select Metric** Select the metric to display in the heatmap. Each Metric has a color gradient bar that maps relative values to colors.

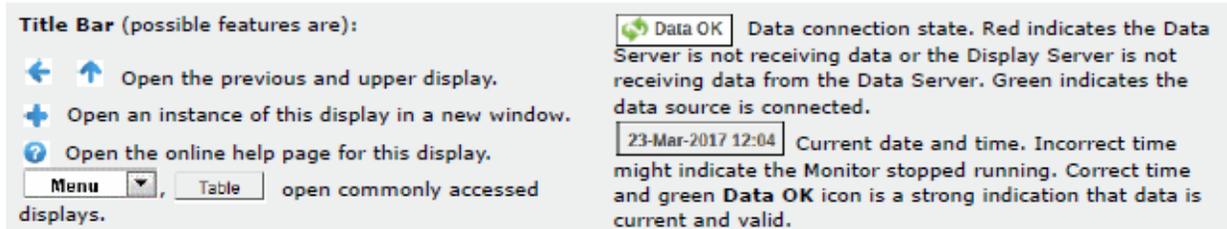
### Applications Summary

Track the performance of all web application modules in a server and view utilization details. The table summarizes the sessions, accesses, cache hit and so forth, for all installed web modules. Each row in the table is a different web application module. The row color for inactive modules is dark red. Select a web application module to view metrics in the trend graph.

Use this data to verify response times of your Web application modules.

Use the available drop-down menus or right-click to filter data shown in the display.





## Fields and Data

This display includes:

- Source** Select the host where the Tomcat Server is running.
- Connection** Select a Tomcat Server from the drop-down menu. This menu is populated by the selected Source.
- Web Module** Select a Web module from the drop-down menu. This menu is populated by the selected Connection. The Web Module you select populates the trend graphs.

### Web Module Summary

<b>Web Module</b>	The name of the Web module.
<b>Sessions Active</b>	The number of currently active client sessions.
<b>Sessions Total</b>	The total number of client sessions since the application was started.
<b>Sessions Expired</b>	The total number of client sessions that expired since the application was started.
<b>Accesses per sec</b>	The number of times pages are accessed, per second.
<b>Accesses Total</b>	The total number of times pages have been accessed since the application was started.
<b>Bytes Rcvd per sec</b>	The number of bytes received per second.
<b>Bytes Rcvd Total</b>	The total number of bytes received since the application was started.
<b>Bytes Sent per sec</b>	The number of bytes sent per second.
<b>Bytes Sent Total</b>	The total number of bytes sent since the application was started.
<b>Cache Hit Rate</b>	The number of times the cache is accessed, per second.
<b>Requests per sec</b>	The number of requests received, per second.
<b>Requests Total</b>	The total number of requests received since the application was started.
<b>Process Time</b>	The average amount of time, in milliseconds, to process requests.
<b>Error Count</b>	The number of errors occurred since the application was started.
<b>appBase</b>	The directory in which Tomcat is installed.

<b>Expired</b>	When checked, this connection is expired due to inactivity.
<b>time_stamp</b>	The date and time this row of data was last updated. Format: <b>MM/DD/YY HH:MM:SS</b> <b>&lt;month&gt;/ &lt;day&gt;/&lt;year&gt;</b> <b>&lt;hours&gt;:&lt;minutes&gt;:&lt;seconds&gt;</b>

**Session/Data/Latency Trends**

Shows metrics for the selected Web module. The Web module can be selected from the **Web Module** drop-down menu or the **Web Modules Summary** table.

<b>Log Scale</b>	Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.
<b>Base at Zero</b>	Use zero as the Y axis minimum for all graph traces.
<b>Time Range</b>	Select a time range from the drop down menu varying from <b>2 Minutes</b> to <b>Last 7 Days</b> , or display <b>All Data</b> . To specify a time range, click Calendar <input type="button" value="..."/> .

By default, the time range end point is the current time. To change the time range end point, click Calendar  and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows   to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu.

Click **Restore to Now** to reset the time range end point to the current time.

<b>Active Sessions</b>	Traces the number of currently active client sessions.
<b>Accesses /sec</b>	Traces the number of times pages are accessed, per second.
<b>Process Time</b>	Traces the average amount of time, in milliseconds, to process requests.

---

## Tomcat Displays - HTML

The Tomcat HTML displays provide extensive visibility into the health and performance of Tomcat application servers and installed web modules. The following Tomcat Views (and their associated displays) can be found under **Components** tab > **Application/Web Servers** > **Tomcat**.

Tomcat has the following displays:

- ["Tomcat Overview - HTML"](#)
- ["Tomcat Servers Heatmap - HTML"](#): Performance metrics for one Tomcat Server, including current and historic performance metrics.
- ["Single Tomcat Server - HTML"](#): Heatmap of performance metrics for all Web modules for one Tomcat Server.
- ["All Tomcat Apps - HTML"](#): Table and trend graphs of performance metrics for Web modules.
- ["Single Tomcat App - HTML"](#): Table and trend graphs of performance metrics for a single Web module.

### Tomcat Overview - HTML

The Tomcat Overview is the top-level display for the Tomcat Solution Package, which provides a good starting point for immediately getting the status of all your Tomcat servers, web modules and connections. You can select the RTView DataServer for which you want to see data and easily view the current data for that DataServer including:

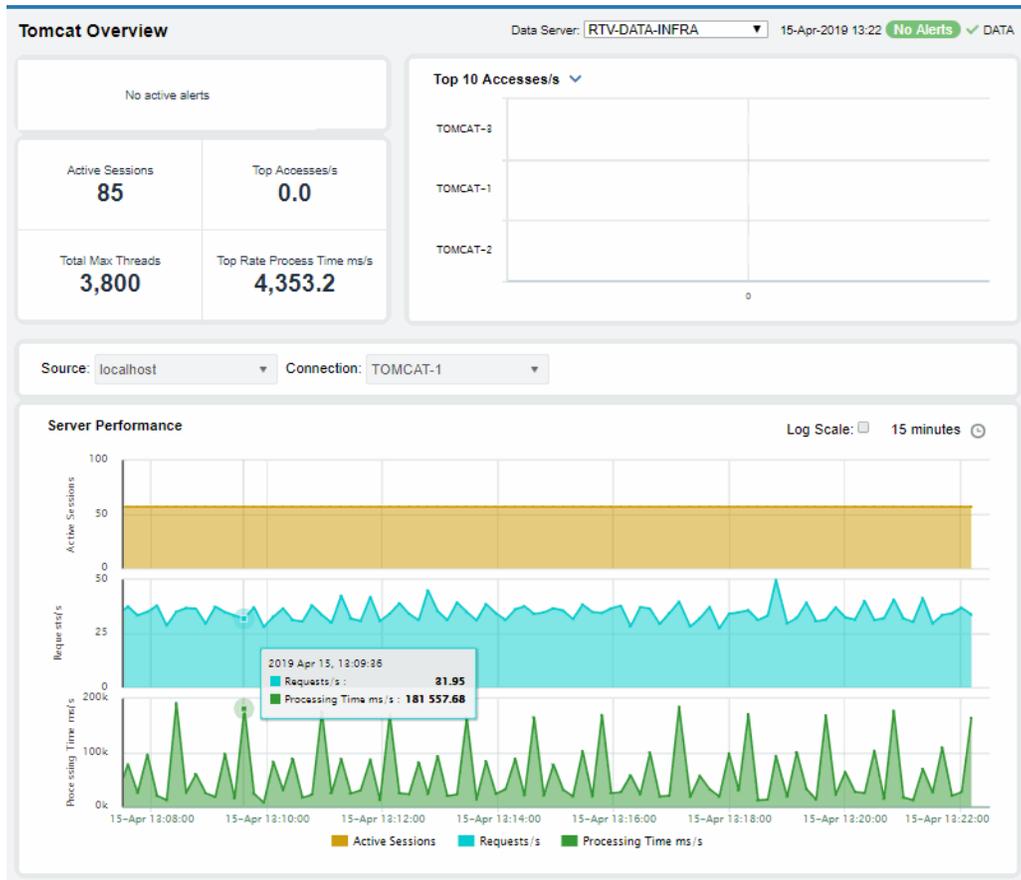
- The total number of active alerts for the selected DataServer, including the total number of critical and warning alerts.
- The greatest number of active sessions, top accesses per second, highest number of connections and the top process rate.
- A visual list of the top 10 servers with the greatest values for **Accesses, Requests, Cache Hit Rate, Process Rate, Sent** and **Received Rate**.

You can hover over each region in the upper half of the Overview to see more detail in a Summary display.

For example, clicking on the alerts in the CRITICAL and WARNING alerts region opens the **Alerts Table by Components** display.

The bottom half of the display provides a trend graph which traces **Active Sessions, Requests** per second and **Processing Time**.

You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.

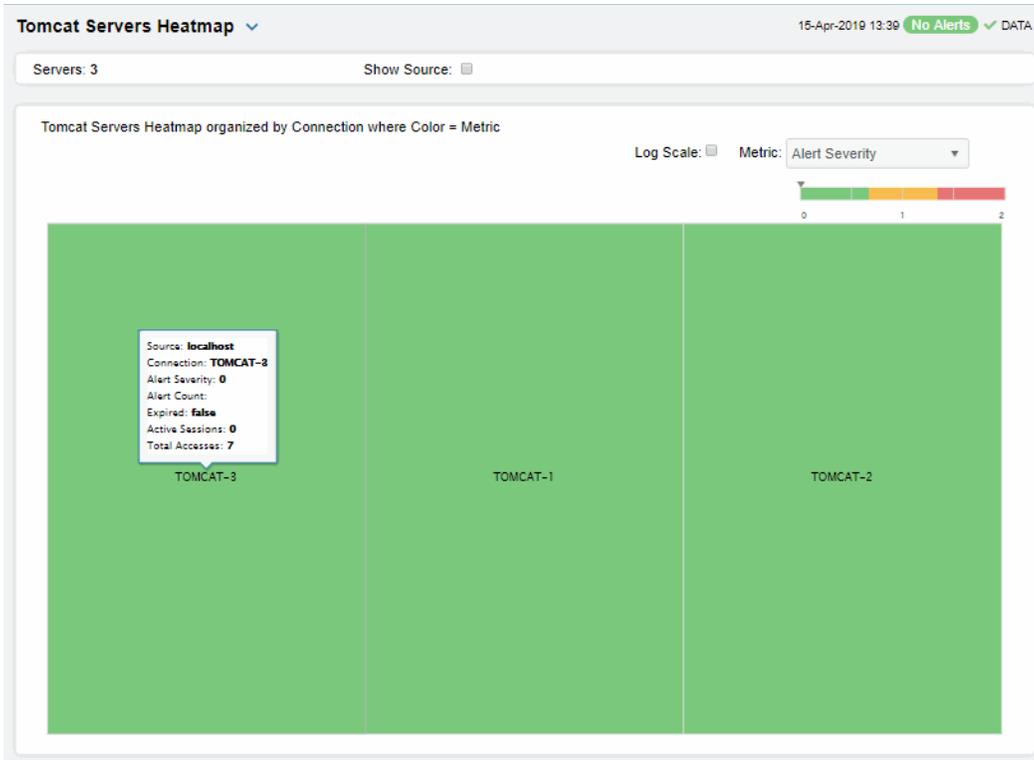


## Tomcat Servers Heatmap - HTML

View performance metrics for all monitored Tomcat Servers. The heatmap organizes Tomcat Web modules by server, and uses color to show the most critical Metric value for each Tomcat connection associated with the selected source. Each rectangle in the heatmap represents a Web module. In this heatmap, the rectangle size is the same for all Web modules. Each Metric (selected from the drop-down menu) has a color gradient bar that maps relative values to colors.

Use this display to see at-a-glance the health of all your web applications. You can select the heatmap color metric from a list including active sessions, access rate, and total access count.

Use the available drop-down menus or right-click to filter data shown in the display. Use the check-boxes  to include or exclude labels in the heatmap. Move your mouse over a rectangle to see additional information. Drill-down and investigate by clicking a rectangle in the heatmap to view details for the selected Web module in the **Application Summary** display.

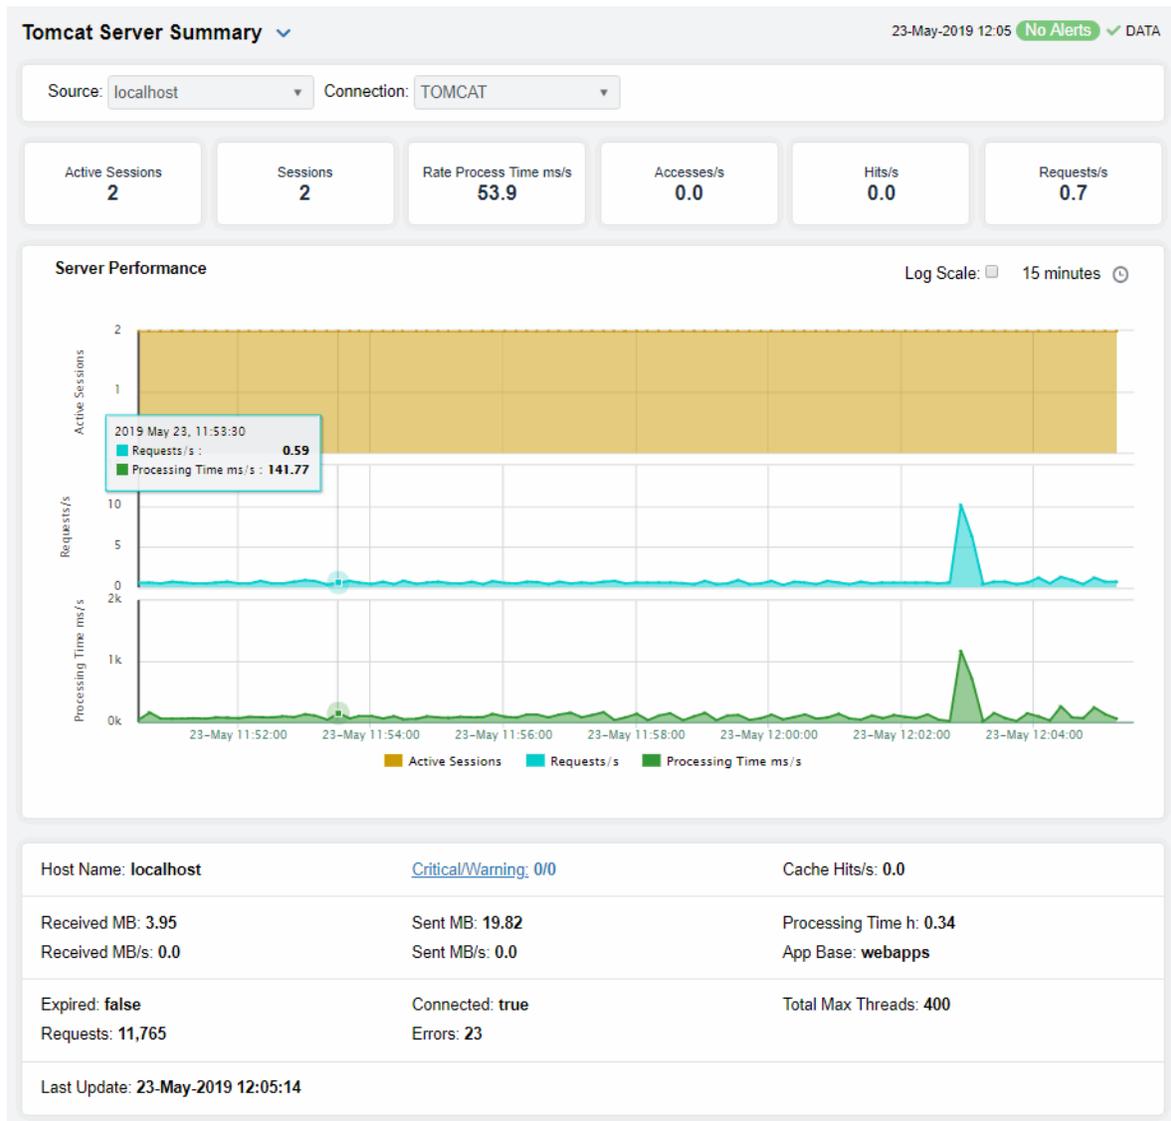


## Single Tomcat Server - HTML

Track utilization and performance metrics for a connection on a Tomcat server. Clicking on the sessions/processing rate information boxes at the top of the display takes you to the **Tomcat Servers Table** display, where you can compare and sort performance values against other Tomcat servers.

The trend graph traces for **Processing Time per second**, **Requests per second** and (number of) **Active Sessions**. You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.

Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display.



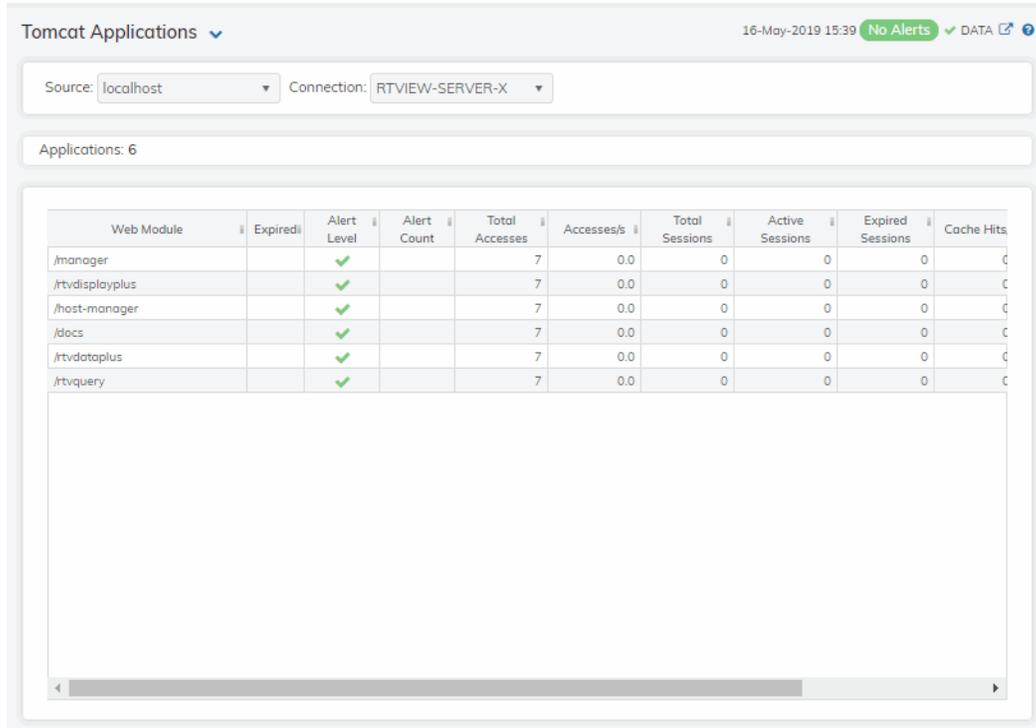
## Tomcat Applications - HTML

Investigate detailed utilization metrics for all Tomcat applications. This display contains all metrics available for Tomcat applications, including the total **Alert Count**, **Accesses/per second** and **Total Sessions**.

Choose a particular **Source** or **All**, and a particular **Connection** or **All**, from the drop-downs. Each row in the table contains data for a particular web module. You can search, filter, sort and choose columns to include by clicking a column header icon (to the right of each column label) and selecting **Filter**, **Sort Ascending**, **Sort Descending** or **Columns**.

To investigate further, double-click a web module to see details in the **Tomcat Application Summary display**.

Or just click a column header to sort. Right-click on a table cell to **Export to Excel** or **Copy Cell Value**.



Tomcat Applications 16-May-2019 15:39 No Alerts DATA

Source: localhost Connection: RTVIEW-SERVER-X

Applications: 6

Web Module	Expired	Alert Level	Alert Count	Total Accesses	Accesses/s	Total Sessions	Active Sessions	Expired Sessions	Cache Hits
/manager		OK		7	0.0	0	0	0	C
/rtvdisplayplus		OK		7	0.0	0	0	0	C
/host-manager		OK		7	0.0	0	0	0	C
/docs		OK		7	0.0	0	0	0	C
/rtvdataplus		OK		7	0.0	0	0	0	C
/rtvquery		OK		7	0.0	0	0	0	C

## All Tomcat Apps - HTML

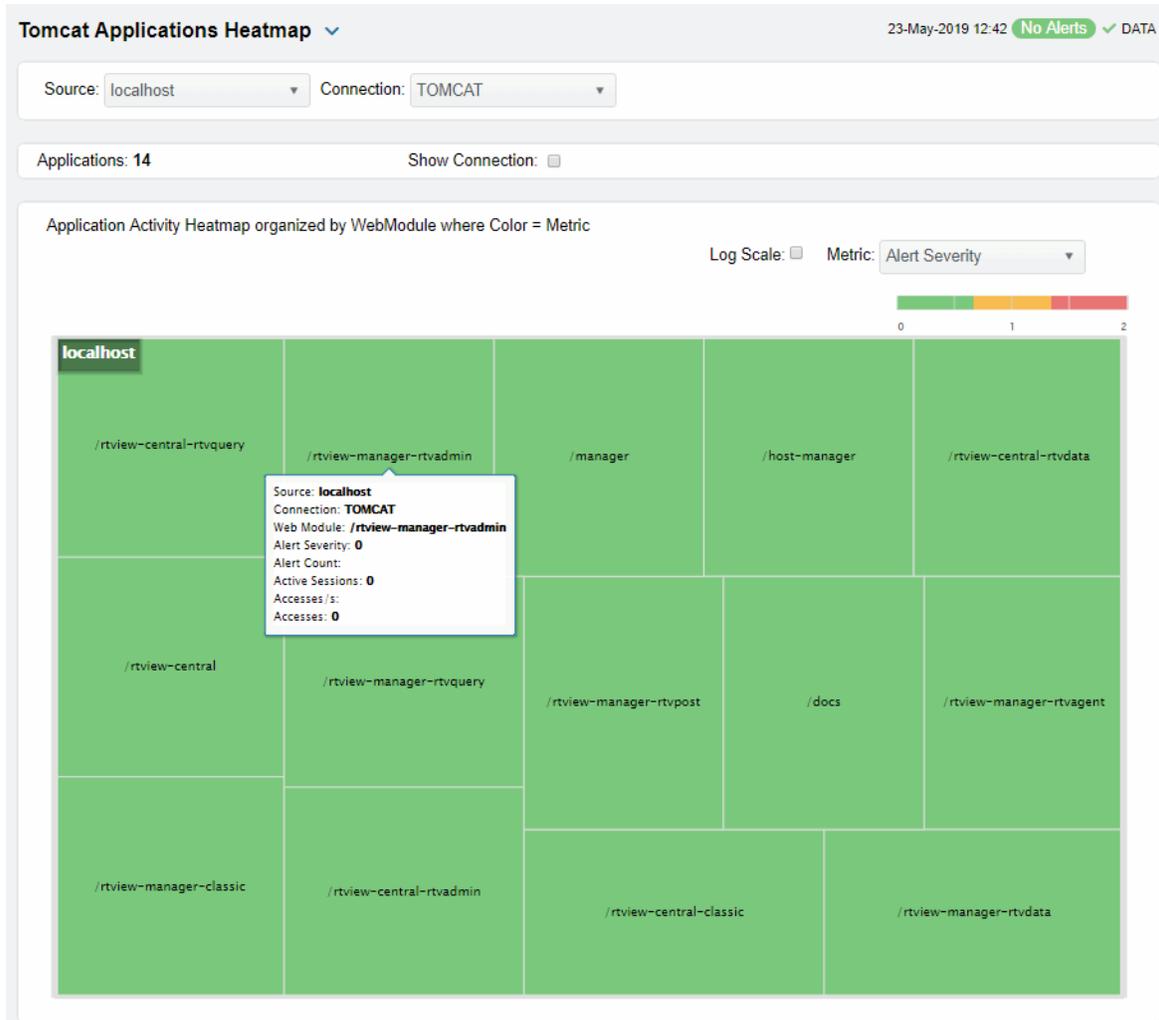
This heatmap allows you to view the status and alerts of Tomcat applications on a particular host or **All** hosts, and a particular connection or **All** connections.

Use the **Metric** drop-down menu to view the **Alert Severity**, **Alert Count**, **Active Sessions**, **Accesses per Second** or (the total number of) **Accesses**.

Each rectangle in the heatmap represents a web module. The rectangle color indicates the most critical alert state. Click on a rectangle to drill-down to the **Tomcat Application Summary display** and view metrics for a particular web module. Toggle between the commonly accessed Table and Heatmap displays by clicking the drop down list on the display title.

Mouse-over rectangles to view more details about host performance and status.

You can view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.



## Single Tomcat App - HTML

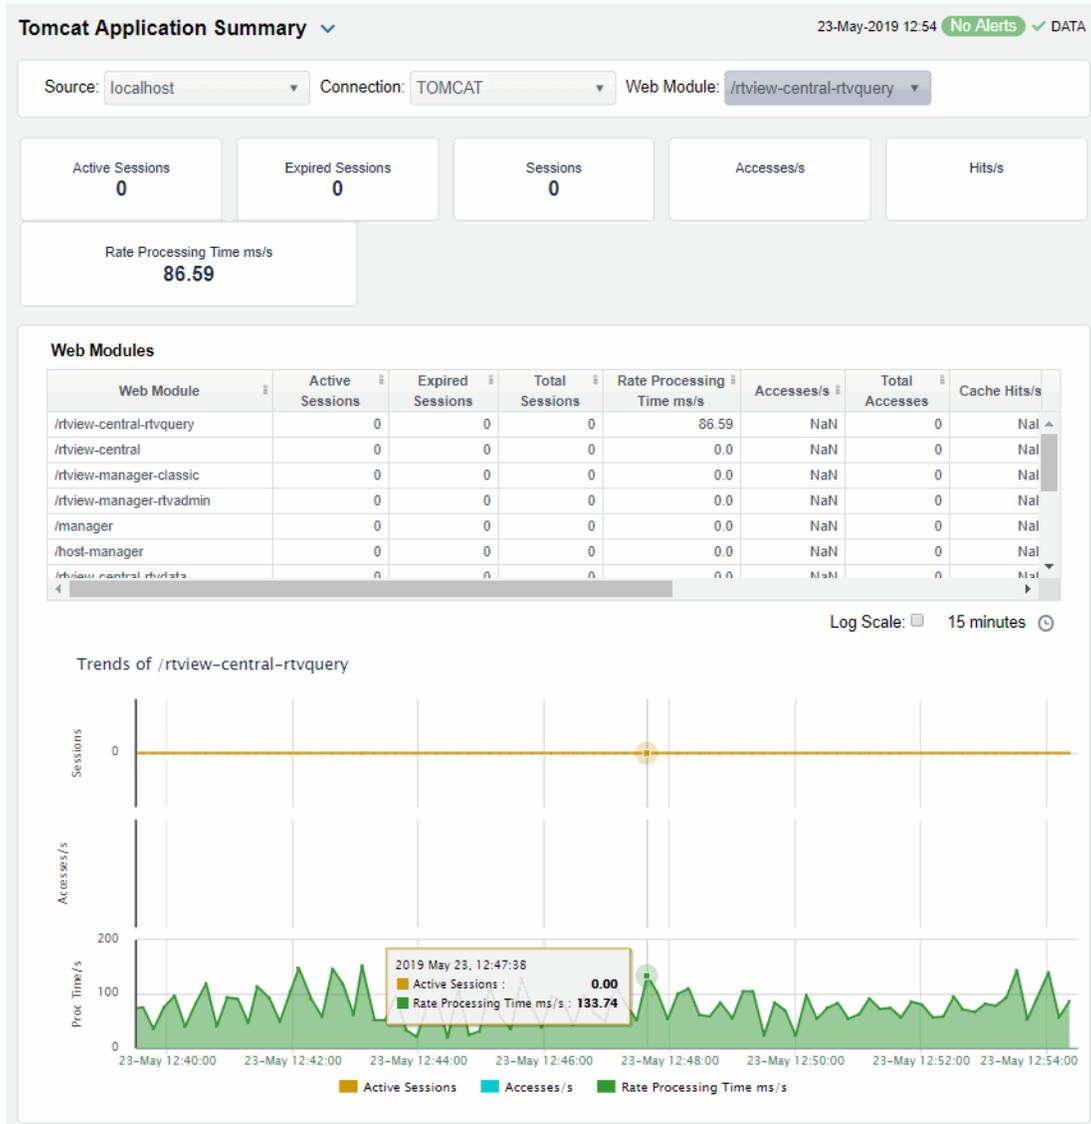
Track utilization and performance metrics for a particular Tomcat web module. Clicking on the sessions/processing rate information boxes at the top of the display takes you to the **Tomcat Servers Table** display, where you can compare and sort performance values against other Tomcat servers.

Use the **Web Modules** table to compare detailed utilization metrics for all web modules. Each row in the table contains data for a particular web module. You can search, filter, sort and choose columns to include by clicking a column header icon (to the right of each column label) and selecting **Filter**, **Sort Ascending**, **Sort Descending** or **Columns**.

Or just click a column header to sort. Right-click on a table cell to **Export to Excel** or **Copy Cell Value**.

The trend graph traces for **Processing Time per second**, **Accesses per second** and (the number of) **Active Sessions**. You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.

Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display.





## APPENDIX A Alert Definitions

This section describes alerts that are available with RTView Enterprise per solution package. This section includes:

- "Amazon Web Services"
- "Apache Kafka"
- "Docker"
- "Microsoft SQL Server"
- "MongoDB"
- "MySQL Database"
- "Node.js"
- "Oracle Coherence"
- "Oracle Database"
- "Oracle WebLogic"
- "RTView Manager and RTView Rules"
- "RTView Host Agent"
- "Solace"
- "TIBCO ActiveMatrix BusinessWorks"
- "TIBCO ActiveSpaces"
- "TIBCO ActiveSpaces (2.x)"
- "TIBCO Adapters"
- "TIBCO BusinessEvents"
- "TIBCO Enterprise Message Service"
- "TIBCO FTL"
- "UX"
- "VMware vCenter"

---

### Amazon Web Services

The following alerts are available for Amazon Web Services. Default settings for warning and alarm thresholds, duration and whether the alert is enabled (true/false) are shown.

Alert	Warning Level	Alarm Level	Duration	Enabled
<b>AcwInstanceCpuHigh</b> Executes a single warning and single alarm if the CPU used exceeds the specified threshold. Index Type: PerInstance Metric: CPUUtilization	70	80	30	FALSE
<b>AcwInstanceDiskReadBytesHigh</b> Executes a single warning and single alarm if the number of bytes read from the disk exceeds the specified threshold. Index Type: PerInstance Metric: DiskReadBytes	10000	20000	30	FALSE
<b>AcwInstanceDiskReadOpsHigh</b> Executes a single warning and single alarm if the number of disk reads exceeds the specified threshold. Index Type: PerInstance Metric: DiskReadOps	100	200	30	FALSE
<b>AcwInstanceDiskWriteBytesHigh</b> Executes a single warning and single alarm if the number of bytes written to the disk exceeds the specified threshold. Index Type: PerInstance Metric: DiskWriteBytes	1000000	2000000	30	FALSE
<b>AcwInstanceDiskWriteOpsHigh</b> Executes a single warning and single alarm if the number of disk writes exceeds the specified threshold. Index Type: PerInstance Metric: DiskWriteOps	100	200	30	FALSE
<b>AcwInstanceNetworkReadBytesHigh</b> Executes a single warning and single alarm if the number of bytes read from the network exceeds the specified threshold. Index Type: PerInstance Metric: NetworkIn	1000000	20000	30	FALSE
<b>AcwInstanceNetworkWriteBytesHigh</b> Executes a single warning and single alarm if the number of bytes written across the network exceeds the specified threshold. Index Type: PerInstance Metric: NetworkOut	10000	20000	30	FALSE

## Apache Kafka

The following alerts are available for Apache Kafka. Default settings for warning and alarm thresholds, duration and whether the alert is enabled (true/false) are shown.

Alert Name	WARN. LEVEL	ALARM LEVEL	DURATION	ENABLED
<b>KafkaBrokerBytesInPerSecHigh</b> The number of incoming bytes per second exceeds the defined threshold for the broker. <b>Index Type(s):</b> PerKafkaServer	1600	2000	30	FALSE
<b>KafkaBrokerBytesOutPerSecHigh</b> The number of outgoing bytes per second exceeds the defined threshold for the broker. <b>Index Type(s):</b> PerKafkaServer	1600	2000	30	FALSE
<b>KafkaBrokerExpired</b> The Kafka Broker is not responding. <b>Index Type(s):</b> PerKafkaServer	NaN	NaN	30	FALSE
<b>KafkaBrokerFetchRequestsPerSecHigh</b> Fetch requests per second exceeds threshold for the broker. <b>Index Type(s):</b> PerKafkaServer	1600	2000	30	FALSE
<b>KafkaBrokerLogFlushLatency95PHigh</b> The current log flush latency exceeds the 95th percentile. <b>Index Type(s):</b> PerKafkaServer	1600	2000	30	FALSE
<b>KafkaBrokerMsgsInPerSecHigh</b> The number of incoming messages per second exceeds the defined threshold for the broker. <b>Index Type(s):</b> PerKafkaServer	1600	2000	30	TRUE
<b>KafkaBrokerNetProcAvgIdlePctHigh</b> The average percent idle for the network processor exceeds the threshold. <b>Index Type(s):</b> PerKafkaServer	.05	.3	30	FALSE
<b>KafkaBrokerProduceRequestsPerSecHigh</b> Produce requests per second exceeds threshold for the broker. <b>Index Type(s):</b> PerKafkaServer	1600	2000	30	FALSE
<b>KafkaBrokerOfflinePartitionCountHigh</b> The number of partitions without an active leader is not zero. <b>Index Type(s):</b> PerKafkaServer	NaN	1	30	TRUE
<b>KafkaBrokerExpired</b> The Kafka Broker is not responding. <b>Index Type(s):</b> PerKafkaServer	NaN	NaN	30	FALSE

<b>KafkaBrokerUncInLderElecsPerSecHigh</b> The available replicas were not in sync during leader election. Data loss has probably occurred. <b>Index Type(s):</b> PerKafkaServer	1600	2000	30	FALSE
<b>KafkaBrokerUnderReplicatedPartnsHigh</b> The number of under-replicated partitions is not zero. <b>Index Type(s):</b> PerKafkaServer	NaN	1	30	FALSE
<b>KafkaClusterLeadersUnbalancedHigh</b> The partition leaders for the cluster are not evenly distributed across the available brokers. <b>Index Type(s):</b> PerKafkaCluster	10	10	30	FALSE
<b>KafkaClusterNoActiveController</b> There is more than one active controller per cluster, which could indicate a split-brain error. <b>Index Type(s):</b> PerKafkaCluster	NaN	NaN	30	FALSE
<b>KafkaClusterPartitionsUnbalancedHigh</b> Partitions supported by the cluster are not evenly distributed across the available brokers. <b>Index Type(s):</b> PerKafkaCluster	10	1	30	FALSE
<b>KafkaClusterSplitBrain</b> One (or more) zookeeper/broker is not acting as part of the main cluster. <b>Index Type(s):</b> PerKafkaCluster	NaN	NaN	30	FALSE
<b>KafkaoConsumerBytesPerSecHigh</b> The consumer message load (bytes per second) exceeds the threshold. <b>Index Type(s):</b> PerKafkaConsumer	1600	2000	30	FALSE
<b>KafkaConsumerExpired</b> The consumer is not responding. <b>Index Type(s):</b> PerKafkaConsumer	NaN	NaN	30	FALSE
<b>KafkaConsumerFetchLatencyHigh</b> The consumer fetch latency exceeds the threshold. <b>Index Type(s):</b> PerKafkaConsumer	1600	2000	30	TRUE
<b>KafkaConsumerFetchRateHigh</b> The consumer is pulling records from Kafka at a slower than expected rate. <b>Index Type(s):</b> PerKafkaConsumer	1600	2000	30	FALSE
<b>KafkaConsumerLagIncreasing</b> The consumer lag rate of change is greater than zero for the specified duration, which could mean that lag is steadily increasing. <b>Index Type(s):</b> PerKafkaConsumer	NaN	NaN	300	FALSE
<b>KafkaConsumerMaxLagHigh</b> The consumer is falling too far behind the producer. <b>Index Type(s):</b> PerKafkaConsumer	1600	2000	30	TRUE

<p><b>KafkaConsumerRecordsConsumedRateHigh</b> The consumer message load (messages per second) exceeds the threshold. <b>Index Type(s):</b> PerKafkaConsumer</p>	1600	2000	30	TRUE
<p><b>KafkaConsumerPartitionStalled</b> The consumer lag delta is not negative and the current offset delta is positive for the defined duration for a topic on a partition, which could mean that new messages are being added to the partition but the consumer is not reading them. <b>Index Type(s):</b> PerKafkaConsumer</p>	NaN	NaN	300	FALSE
<p><b>KafkaConsumerRecordsConsumedRateHigh</b> Consumer records consumed rate exceeds threshold. <b>Index Type(s):</b> PerKafkaConsumer</p>	1600	2000	300	FALSE
<p><b>KafkaConsumerSlow</b> This alert is triggered for a topic when consumer lag delta is not negative and the current offset delta is positive for the specified duration, which could mean that the consumer is slow in reading messages. <b>Index Type(s):</b> PerKafkaConsumer</p>	NaN	NaN	300	FALSE
<p><b>KafkaProducerExpired</b> The producer is not responding. <b>Index Type(s):</b> PerKafkaProducer</p>	NaN	NaN	30	FALSE
<p><b>KafkaProducerIncomingByteRateHigh</b> The producer's incoming byte rate exceeds the threshold. <b>Index Type(s):</b> PerKafkaProducer</p>	1600	2000	30	TRUE
<p><b>KafkaProducerIoWaitTimeMSHigh</b> The producer is waiting for IO longer than expected (on average). <b>Index Type(s):</b> PerKafkaProducer</p>	1600	2000	30	FALSE
<p><b>KafkaProducerOutgoingByteRateHigh</b> The producer output byte rate exceeds the threshold. <b>Index Type(s):</b> PerKafkaProducer</p>	1600	2000	30	TRUE
<p><b>KafkaProducerRecordSendRateHigh</b> The producer record send rate exceeds the threshold. <b>Index Type(s):</b> PerKafkaProducer</p>	1600	2000	30	TRUE
<p><b>KafkaProducerRequestLatencyHigh</b> The producer request latency exceeds the threshold. <b>Index Type(s):</b> PerKafkaProducer</p>	1600	2000	30	TRUE
<p><b>KafkaProducerRequestRateHigh</b> The producers request rate exceeds the threshold. <b>Index Type(s):</b> PerKafkaProducer</p>	1600	2000	30	TRUE

<b>KafkaProducerResponseRateHigh</b> The producer response rate exceeds the threshold. <b>Index Type(s):</b> PerKafkaProducer	1600	2000	30	TRUE
<b>KafkaZookeeperAvgLatencyHigh</b> The average time for the zookeeper to respond to a request exceeds the threshold. <b>Index Type(s):</b> PerKafkaZookeeper	1600	2000	30	TRUE
<b>KafkaZookeeperCpuPercentHigh</b> The CPU percentage reported by the JVM is above the limits defined for that zookeeper. <b>Index Type(s):</b> PerKafkaZookeeper	50	75	30	TRUE
<b>KafkaZookeeperExpired</b> The zookeeper is not responding. <b>Index Type(s):</b> PerKafkaZookeeper	NaN	NaN	30	FALSE
<b>KafkaZookeeperMemoryUsedPercentHigh</b> The percentage of heap memory used relative to the maximum heap available is above the limits defined for that Zookeeper. <b>Index Type(s):</b> PerKafkaZookeeper	50	75	30	FALSE
<b>KafkaZookeeperNumAliveConnsHigh</b> The total number of connections to a given zookeeper exceeds the threshold. <b>Index Type(s):</b> PerKafkaZookeeper	1600	2000	30	TRUE
<b>KafkaZookeeperOutstandingReqsHigh</b> Clients are making requests faster than the zookeeper can process them. <b>Index Type(s):</b> PerKafkaZookeeper	1600	2000	30	FALSE
<b>KafkaZookeeperRatePktsRcvdHigh</b> The rate that the zookeeper is receiving packets exceeds the threshold. <b>Index Type(s):</b> PerKafkaZookeeper	1600	2000	30	TRUE
<b>KafkaZookeeperRatePktsSentHigh</b> The rate that the zookeeper is sending packets exceeds the threshold. <b>Index Type(s):</b> PerKafkaZookeeper	1600	2000	30	TRUE

---

## Docker

The following alerts are available for Docker. Default settings for warning and alarm thresholds, duration and whether the alert is enabled (true/false) are shown.

Alert Name	WARN. LEVEL	ALARM LEVEL	DURATION	ENABLED
<b>DocContainerCpuUsageHigh</b> A Docker Container's CPU usage is above the defined threshold. <b>Index Type(s):</b> PerContainer <b>Metric:</b> cpu.usage	24	50	30	FALSE
<b>DocContainerExpired</b> A Docker Container has expired. <b>Index Type(s):</b> PerContainer <b>Metric:</b> Expired	NaN	NaN	30	FALSE
<b>DocContainerNetBytesInHigh</b> A Docker Container's incoming network data rate is above the defined thresholds. <b>Index Type(s):</b> PerContainer <b>Metric:</b> net.rxbytes.avg	750000	1000000	30	FALSE
<b>DocContainerNetBytesOutHigh</b> A Docker Container's outgoing network data rate is above the defined thresholds. <b>Index Type(s):</b> PerContainer <b>Metric:</b> net.txbytes.avg	750000	1000000	30	FALSE
<b>DocEngineCpuUsageHigh</b> A Docker Engine's CPU usage is above the defined thresholds. <b>Index Type(s):</b> PerEngine <b>Metric:</b> cpu.usage	50	75	30	TRUE
<b>DocEngineExpired</b> A Docker Engine has expired. <b>Index Type(s):</b> PerEngine <b>Metric:</b> Expired	NaN	NaN	30	FALSE
<b>DocEngineNetBytesInHigh</b> A Docker Engine's incoming network data rate is above the defined thresholds. <b>Index Type(s):</b> PerEngine <b>Metric:</b> net.rxbytes.avg	750000	1000000	30	TRUE
<b>DocEngineNetBytesOutHigh</b> A Docker Engine's outgoing network data rate is above the defined thresholds. <b>Index Type(s):</b> PerEngine <b>Metric:</b> net.txbytes.avg	750000	1000000	30	TRUE

---

## Microsoft SQL Server

The following alerts are available for Microsoft SQL Server. Default settings for warning and alarm thresholds, duration and whether the alert is enabled (true/false) are shown.

Alert Name	WARN. LEVEL	ALARM LEVEL	DURATION	ENABLED
<b>MssqlInstanceDeadlocksDetected</b> The number of current deadlocks has exceeded its threshold. <b>Index Type(s):</b> PerServer <b>Metric:</b> DeltaNumber of Deadlocks	1	2	0	TRUE
<b>MssqlInstanceLatchWaitsHigh</b> The number of current latch waits has exceeded its threshold. <b>Index Type(s):</b> PerServer <b>Metric:</b> DeltaLatch Waits	15	30	0	TRUE
<b>MssqlInstanceLockWaitsHigh</b> The amount of seconds on lock waits has exceeded its threshold. <b>Index Type(s):</b> PerServer <b>Metric:</b> DeltaLock Waits	15	30	0	TRUE
<b>MssqlInstancePacketErrorsDetected</b> The amount of current packet errors has exceeded its threshold. <b>Index Type(s):</b> PerServer <b>Metric:</b> DeltaPacket Errors	1	2	0	TRUE
<b>MssqlInstanceSqlCpuUsedHigh</b> The percentage of CPU utilization on SQL processing has exceeded its threshold. <b>Index Type(s):</b> PerServer <b>Metric:</b> CPU Util	10	15	0	TRUE
<b>MssqlInstanceUsedMemoryHigh</b> The percentage of memory used by the SQL Server has exceeded its threshold. <b>Index Type(s):</b> PerServer <b>Metric:</b> Memory In Use (%)	65	85	0	TRUE

---

## MongoDB

The following alerts are available for MongoDB. Default settings for warning and alarm thresholds, duration and whether the alert is enabled (true/false) are shown.

Alert Name	WARN. LEVEL	ALARM LEVEL	DURATION	ENABLED
<b>MongoCollectionExpired</b> A collection was not able to be contacted for longer than the normal expiration window. <b>Index Type(s):</b> PerCollection <b>Metric:</b> Expired	NaN	NaN	30	FALSE
<b>MongoCollectionNumObjectsHigh</b> The number of objects for the collection exceeds a given threshold. <b>Index Type(s):</b> PerCollection <b>Metric:</b> numberOfObjects	1600	2000	30	FALSE
<b>MongoDatabaseDataSizeHigh</b> The database size for the database exceeds a given threshold. <b>Index Type(s):</b> PerDatabase <b>Metric:</b> dataSize	80000	100000	30	FALSE
<b>MongDatabaseExpired</b> The database was not able to be contacted for longer than the normal expiration window. <b>Index Type(s):</b> PerDatabase <b>Metric:</b> Expired	NaN	NaN	30	FALSE
<b>MongoInstanceExpired</b> The instance was not able to be contacted for longer than the normal expiration window. <b>Index Type(s):</b> PerInstance <b>Metric:</b> Expired	60	80	30	FALSE
<b>MongoInstanceNotConnected</b> The instance was not able to be contacted for longer than the normal expiration window. <b>Index Type(s):</b> PerInstance <b>Metric:</b> connectionStatus	NaN	NaN	30	FALSE
<b>MongoInstanceOpenCursorsHigh</b> The number of Open Cursors for the Instance exceeds a given threshold. <b>Index Type(s):</b> PerInstance <b>Metric:</b> openCursors	160	200	30	TRUE

---

## MySQL Database

The following alerts are available for MySQL Database. Default settings for warning and alarm thresholds, duration and whether the alert is enabled (true/false) are shown.

Alert	Warning Level	Alarm Level	Duration	Enabled
<p><b>MysqlBytesReceivedHigh</b> Executes a single warning and a single alarm if the amount of kilobytes received exceeds the specified threshold. Index Type: PerServer Metric: Received</p>	5	10	na	FALSE
<p><b>MysqlBytesSentHigh</b> Executes a single warning and a single alarm if the amount of kilobytes sent exceeds the specified threshold. Index Type: PerServer Metric: Sent</p>	5	10	na	FALSE
<p><b>MysqlDelayedWritesHigh</b> Executes a single warning and a single alarm if the number of delayed writes exceeds the specified threshold. This alert only applies to previous versions to MySQL 5.7 as delayed inserts are not supported in later versions. Index Type: PerServer Metric: Delayed Writes</p>	1	2	na	FALSE
<p><b>MysqlLocksWaitedHigh</b> Executes a single warning and a single alarm if the number of times that requests for a table lock requires a wait before being granted exceeds the specified threshold. Index Type: PerServer Metric: Table_locks_waited</p>	1	2	na	FALSE
<p><b>MysqlQcacheLowMemPrunesHigh</b> Executes a single warning and a single alarm if the number of queries deleted from the query cache because of low memory exceeds the specified threshold. Index Type: PerServer Metric: Qcache_lowmem_prunes</p>	1	2	na	FALSE

<b>MySQLSlowQueriesHigh</b> Executes a single warning and a single alarm if the number of queries that exceed the number of seconds specified for <b>long_query_time</b> exceeds the specified threshold. Index Type: PerServer Metric: Slow Queries	1	2	na	FALSE
<b>MySQLSlowThreadsHigh</b> Executes a single warning and a single alarm if the number of threads that exceed the number of seconds specified for <b>slow_launch_time</b> to create exceeds the specified threshold. Index Type: PerServer Metric: Slow_launch_threads	1	2	na	FALSE

## Node.js

The following alerts are available for Node.js. Default settings for warning and alarm thresholds, duration and whether the alert is enabled (true/false) are shown.

Alert Name	WARN. LEVEL	ALARM LEVEL	DURATION	ENABLED
<b>NodeMasterCpuUsageHigh</b> A master node's CPU usage is above the defined thresholds. <b>Index Type(s):</b> PerConnection <b>Metric:</b> Node Master - CPU %	30	50	30	FALSE
<b>NodeMasterExpired</b> A master node has expired. <b>Index Type(s):</b> PerConnection <b>Metric:</b> Node Master - Expired	NaN	NaN	30	FALSE
<b>NodeMasterRequestRateHigh</b> The request rate of a master node is above the defined thresholds. <b>Index Type(s):</b> PerConnection <b>Metric:</b> Node Requests - Requests Per Second	1600	2000	30	FALSE
<b>NodeMasterResponseTimeHigh</b> The response time of a URL is above the defined thresholds. <b>Index Type(s):</b> PerConnection <b>Metric:</b> Node Requests - Avg Response Time	5	10	30	FALSE

<b>NodeProcessCpuUsageHigh</b> A worker node's CPU usage is above the defined thresholds. <b>Index Type(s):</b> PerConnection <b>Metric:</b> Node Processes - CPU Used %	5	50	30	TRUE
<b>NodeProcessExpired</b> A worker node has expired. <b>Index Type(s):</b> PerConnection <b>Metric:</b> Node Processes - Expired	NaN	NaN	30	FALSE
<b>NodeProcessMemUsageHigh</b> A master node's memory usage has exceeded the defined limits. <b>Index Type(s):</b> PerConnection <b>Metric:</b> Node Processes - Memory Used %	90	95	30	TRUE

---

## Oracle Coherence

The following alerts are available with both the solution package and standalone versions for Oracle Coherence.

### OcAvailableMemoryLowCluster

A single alert is executed if the average percent memory used over max memory of all nodes in the cluster exceeds the specified thresholds.

### OcAvailableMemoryLowNode

For each node in the cluster, an alert is executed if the percent memory used over max memory available for that node exceeds the specified thresholds.

### OcAvailableMemoryLowNodeSpike

For each node in the cluster, an alert is executed if the percent memory used exceeds the specified threshold for the percent above average memory used in the previous 24 hours. For example, if the threshold is set to 50% of total memory used, and the average memory consumption on a particular node for the previous 24 hours is 40%, an alert will be executed if current memory usage exceeds 60% of the total.

NOTE: The 24 hour time span (86400 seconds) is controlled by the \$AVERAGE\_MEMORY\_TIME\_WINDOW substitution.

The warning default setting is **115** (percent) of the previous 24 hours and the alarm default setting is **125** (percent) of the previous 24 hours.

By default the alert is disabled.

### OcBadCommunicationCluster

A single alert is executed if the average communication failure rate of all nodes in the cluster exceeds the specified thresholds.

### OcBadCommunicationNode

For each node in the cluster, an alert is executed if the communication failure rate for that node exceeds the specified thresholds.

### **OcBadCommunicationNodesInTimeRange**

Executes a single warning and a single alert if the percentage of nodes in a cluster exceeds the specified threshold for the BadCommunicationNode alert within a time range specified.

To specify the time range, modify the \$BAD\_COMMUNICATION\_NODES\_TIME\_RANGE substitution.

The default time range setting is 5 minutes (300 seconds), the warning default setting is **40** (percent) and the alarm default setting is **50** (percent).

By default the alert is enabled.

### **OcCacheHitPercentageLow**

This alert is executed when the current **Hit%** (total current hits/total current gets) is below the specified threshold for a sampling period and the specified cache(s).

### **OcCacheQueueSizeHigh**

A single alert is executed when the CacheQueueSize for all nodes in the cluster exceeds the specified thresholds. By default the alert is disabled with the following default settings: Warning is **100** (total objects), Alarm is **200** (total objects) and Duration is **60** (total objects).

### **OcCacheRateCacheMissesHigh**

Executes when the Misses per second exceed the specified threshold and duration. The rate is for a given tier of a cache for a given service in a cluster. The tier can be front, where appropriate, or back. Caches and services are named, and clusters are represented by their named monitoring connection. This alert has PerCluster, PerService, PerCache and overrides. This alert appears in the Other Category when triggered.

This is a Key Metrics alert that is available with the RTView Enterprise Monitor when the Oracle Coherence Monitor is installed.

By default the alert is disabled with the following default settings: Warning is **1000**, Alarm is **2000** and Duration is **0** (seconds). Before enabling this alert, you **MUST** change the default settings to values that are suitable for your environment.

### **OcCacheRateStoreReadsHigh**

Executes when the cache StoreReads rate per second exceeds the specified thresholds and durations. The rate is for a given tier of a cache for a given service in a cluster. The tier can be front, where appropriate, or back. Caches and services are named, and clusters are represented by their named monitoring connection. This alert has PerCluster, PerService, PerCache and overrides. This alert appears in the Other Category when triggered.

This is a Key Metrics alert that is available with the RTView Enterprise Monitor when the Oracle Coherence Monitor is installed.

By default the alert is disabled with the following default settings: Warning is **1000**, Alarm is **5000** and Duration is **0** (seconds). Before enabling this alert, you **MUST** change the default settings to values that are suitable for your environment.

### **OcCacheRateStoreWritesHigh**

Executes when the cache StoreWrites rate per second exceeds the specified thresholds and durations. The rate is for a given tier of a cache for a given service in a cluster. The tier can be front, where appropriate, or back. Caches and services are named, and clusters are represented by their named monitoring connection. This alert has PerCluster, PerService, PerCache and overrides. This alert appears in the Other Category when triggered.

This is a Key Metrics alert that is available with the RTView Enterprise Monitor when the Oracle Coherence Monitor is installed.

By default the alert is disabled with the following default settings: Warning is **1000**, Alarm is **5000** and Duration is **0** (seconds). Before enabling this alert, you **MUST** change the default settings to values that are suitable for your environment.

### **OcCacheRateTotalGetsHigh**

Executes when the cache total gets rate per second exceeds the specified thresholds and durations. The rate is for a given tier of a cache for a given service in a cluster. The tier can be front, where appropriate, or back. Caches and services are named, and clusters are represented by their named monitoring connection. This alert has PerCluster, PerService, PerCache and overrides. This alert appears in the Other Category when triggered.

This is a Key Metrics alert that is available with the RTView Enterprise Monitor when the Oracle Coherence Monitor is installed.

By default the alert is disabled with the following default settings: Warning is **1000**, Alarm is **5000** and Duration is **0** (seconds). Before enabling this alert, you **MUST** change the default settings to values that are suitable for your environment.

### **OcCacheRateTotalPutsHigh**

Executes when the cache DeltaTotalPuts rate per second exceeds the specified thresholds and durations. The rate is for a given tier of a cache for a given service in a cluster. The tier can be front, where appropriate, or back. Caches and services are named, and clusters are represented by their named monitoring connection. This alert has PerCluster, PerService, PerCache and overrides. This alert appears in the Other Category when triggered.

This is a Key Metrics alert that is available with the RTView Enterprise Monitor when the Oracle Coherence Monitor is installed.

By default the alert is disabled with the following default settings: Warning is **1000**, Alarm is **5000** and Duration is **0** (seconds). Before enabling this alert, you **MUST** change the default settings to values that are suitable for your environment.

### **OCCacheSizeHigh**

Executes when the number of objects in a cache exceeds the specified threshold. By default the alert is disabled with the following default settings: Warning is **1000** (count), Alarm is **5000** (count) and Duration is **60** (seconds).

This is a Key Metrics alert that is available with the RTView Enterprise Monitor when the Oracle Coherence Monitor is installed.

NOTE: If you want to know when the size of a specific cache exceeds specific thresholds, it might be preferable to use the **Per Cache** or **Per Storage Class** override settings, allowing you set specific thresholds for specific caches.

### **OCCacheSizeLow**

Executes when the number of objects in a cache goes below the specified threshold. By default the alert is disabled with the following default settings: Warning is **1000** (count), Alarm is **5000** (count) and Duration is **60** (seconds).

NOTE: If you want to know when the size of a specific cache goes below specific thresholds, it might be preferable to use the **Per Cache** or **Per Storage Class** override settings, allowing you set specific thresholds for specific caches.

### **OcCapacityLimitAllCaches**

An alert is executed if the percent cache used over cache capacity for any cache in the cluster exceeds the specified thresholds. There is one highWarning and one highAlert threshold. For example, if there are 3 caches in a cluster, where:

**cache1 val = 95**

**cache2 val = 100**

**cache3 val = 70**

and the CapacityLimitAllCaches highWarning is **80** and highAlert is **90**, one high alert is executed.

### **OcCapacityLimitCache**

Executes when the average CPU usage for the cluster / storage class exceeds the specified thresholds and durations. This alert has a per cluster and a per (cluster) storage class override. This alert appears in the Other Category when executed.

This is a Key Metrics alert that is available with the RTView Enterprise Monitor when the Oracle Coherence Monitor is installed.

By default the alert is disabled with the following default settings: Warning is **95** (percent), Alarm is **95** (percent) and Duration is **60** (seconds).

### **OcClusterNodesRcvdFailureRateHigh**

Executes when the average network/packet received failure rate for the cluster/storage class exceeds the specified thresholds and durations. The metrics are averaged across all nodes of a storage class in a cluster.

This is a Key Metrics alert that is available with the RTView Enterprise Monitor when the Oracle Coherence Monitor is installed.

This alert has a per cluster and a per (cluster) storage class override. Note that this alert appears in the Network Category when executed.

By default the alert is disabled with the following default settings: Warning is **95** (percent), Alarm is **95** (percent) and Duration is **60** (seconds).

### **OcClusterNodesSentFailureRateHigh**

Executes when the average network/packet sent failure rate for the cluster / storage class exceeds the specified thresholds and durations. The metrics are averaged across all nodes of a storage class in a cluster.

This is a Key Metrics alert that is available with the RTView Enterprise Monitor when the Oracle Coherence Monitor is installed.

This alert has a per cluster and a per (cluster) storage class override. Note that this alert appears in the Memory Category when executed.

By default the alert is disabled with the following default settings: Warning is **95** (percent), Alarm is **95** (percent) and Duration is **60** (seconds).

### **OcDepartedNode**

For each node in the cluster, an alert is executed if the time a node is absent from the cluster exceeds the specified thresholds. When the departed node rejoins the cluster, the alert is cleared.

### **OcDepartedNodesPercentage**

This scalar alert executes a single warning and a single alert if the percentage of nodes departed from the cluster exceeds the specified thresholds within the specified time periods. The percentage is measured against the total number of nodes in the cluster, including both running and departed nodes.

The time period is set in the **rtview.properties** file using the `$NODES_DEPARTED_TIME_WINDOW` substitution. The time period can also be overridden using the command line interface. For example, the following sets a time window of 300 seconds:

```
-sub:$NODES_DEPARTED_TIME_WINDOW:300
```

The time period default setting is **600** (10 minutes), the warning default setting is **90** (percent) and the alarm default setting is **95** (percent).

By default the alert is disabled.

### **OcEndangeredAllCaches**

This alert is executed if the StatusHA for the cache service is `NODE_SAFE` (high warning) or `ENDANGERED` (high alert).

### **OcEndangeredCache**

For each node in the cluster, an alert is executed if the StatusHA value is `ENDANGERED`. By default the alert is disabled.

### **OcExtendConnectionByteBacklogHigh**

This limits alert executes a single warning and a single alert if the `OutgoingByteBacklog` for a Proxy Extend Connection exceeds the specified thresholds. By default the alert is disabled with the following default settings: Warning is **1000** (bytes), Alert is **5000** (bytes).

### **OcHATargetFailed**

This alert executes when the distributed service target status (`HATarget`) is not met. The `HATarget` value is determined using the `PartitionAssignment` MBean in Coherence Versions 12 and above. In prior Coherence versions, the default value of `MACHINE_SAFE` is used. The default value can be overridden by setting the substitution variable `$ocmDefaultHATarget` to the desired value.

### **OcHighGCDutyCycleNode**

This scalar alert executes a single warning and a single alert if a node exceeds the specified duty cycle threshold (the percent of time spent in Garbage Collection).

By default the alert is enabled with the following default settings: Warning is **10** (percent), Alarm is **20** (percent) and Duration is **10** seconds.

### **OcHighPendingRequestNode**

A single alert is executed if the `RequestPendingCount` amount exceeds the specified threshold. This alert allows for setting the warning level, alarm level and duration.

By default the alert is disabled.

### **OcHighTaskBacklogNode**

A single warning and a single alert are executed if the number of backlogged tasks exceeds the specified user threshold. This alert allows for setting the warning level, alarm level and duration.

The default setting executes a warning if the number of backlogged tasks exceeds **10**, and executes an alert if the number of backlogged tasks exceeds **20**.

By default the alert is disabled.

### **OcHighThreadAbandonedNode**

A single alert is executed if the Coherence Thread Abandoned Count amount exceeds the specified threshold. This alert allows for setting the warning level, alarm level and duration.

The default setting executes a warning and an alert if the Thread Abandoned Count amount exceeds **0**. The default duration setting is **60**.

By default the alert is enabled.

### **OcJmxProcessingTime**

This alert is executed if the sum of time for JMX queries and all data processing functions exceeds the specified threshold for the **jmxsampleperiod** property. By default the alert is disabled with the following default settings: Warning is **80** (percent), Alarm is **90** (percent) and Duration is **0** (seconds).

NOTE: The OcJmxProcessingTime alert does not support overrides. For that alert the Override Count is displayed as **-1**.

### **OcLongGCDurationNode**

A single warning and a single alert are executed if any of the last garbage collection times exceed the specified duration.

The default setting executes a warning if the duration exceeds 1 second, and executes an alert if the duration exceeds 2 seconds.

It is possible for GC times to exceed the specified duration and NOT execute an alert. This is possible if it occurs between the alert duration time and an alert condition time.

For example, if your alert duration is 60 seconds, and there is also an alert condition set at 27 seconds into that 60 seconds, the following scenarios could occur (where XX:XX:XX is Hours:Minutes:Seconds):

#### **Scenario 1:**

12:00:00 GC amount is below the specified threshold. No alert executed.

12:00:27 GC amount exceeds the specified threshold. Alert ignored for now.

12:01:00 GC amount is below the specified threshold. No alert executed.

#### **Scenario 2:**

12:00:00 GC amount is below the specified threshold. No alert executed.

12:00:27 GC amount exceeds the specified threshold. Alert ignored for now.

12:01:00 GC amount remains above the specified threshold. Alert executed.

By default the alert is enabled.

### **OcLowClientNodeCount**

This alert executes if the total number of nodes being monitored, including storage enabled nodes, client nodes, and management (JMX) nodes, exceeds the specified threshold. When the count returns to above to above the threshold (departed nodes rejoin the cluster), the alert is cleared.

By default the alert is disabled.

### **OcLowStorageNodeCount**

This alert executes if the total number of storage nodes in the cluster exceeds the specified threshold. When the count returns to above to above the threshold (departed nodes rejoin the cluster), the alert is cleared.

By default the alert is disabled.

### **OcLowTotalNodeCount**

This alert executes if the total number of client nodes being monitored exceeds the specified threshold. When the count returns to above to above the threshold (departed nodes rejoin the cluster), the alert is cleared.

By default the alert is disabled.

### **OcMemoryUsedPercentageAfterGC**

This alert is executed if the percent of memory used on a node after garbage collection exceeds the specified threshold. By default the alert is disabled with the following default settings: Warning is **70** (percent), Alarm is **80** (percent) and Duration is **30** (seconds).

### **OcNodeSafeCache**

For each node in the cluster, an alert is executed if the StatusHA value is **NODE-SAFE**. By default the alert is disabled.

### **OcNoJmxConnection**

This alert is executed if a JMX connection remains disconnected after a specified duration of time. The default duration of time is **60** seconds. By default, this alert is enabled.

### **OcObjectCountDeltaUpCache**

This tabular alert executes a single warning and a single alert for each cache in the cluster if the cache object count delta increases and exceeds the specified threshold. In addition to setting the warning and alarm levels, this alert also allows for setting the duration for each cache.

When this alert is selected in the Active Alert Table, the Per Cache Alert Setting box is displayed (rather than the scalar alert box).

By default the alert is disabled.

### **OcObjectCountDeltaDownCache**

This tabular alert executes a single warning and a single alert for each cache in the cluster where the cache object count delta decreases and exceeds the specified threshold. In addition to setting the warning and alarm levels, this alert also allows for setting the duration for each cache.

When this alert is selected in the Active Alert Table, the Per Cache Alert Setting box is displayed (rather than the scalar alert box).

By default the alert is disabled.

### **OcProxyNodeByteBacklogHigh**

This limits alert executes a single warning and a single alert if the `OutgoingByteBacklog` for a Proxy Node exceeds the specified threshold. This is often indicates overloaded capacity on an individual proxy node. By default the alert is disabled with the following default settings: Warning is **100** (bytes), Alert is **50** (bytes).

### **OcSendQueueSize**

For each node in the cluster, an alert is executed if the Send Queue for that node exceeds the specified thresholds. By default the alert is disabled with the following default settings: Warning is **100** (seconds), Alarm is **200** (seconds) and Duration is **60** (seconds).

### **OcStoreFailure**

This alert is executed if the number of StoreFailures exceeds the specified threshold. By default the alert is disabled with the following default settings: Warning is **1** (second), Alarm is **10** (seconds) and Duration is **30** (seconds).

### **OcStoreReadMillisHigh**

This alert is executed if the current average read per millisecond (total current `StoreReadMillis`/total current `StoreReads`) exceeds the specified threshold for a sampling period and the specified cache(s).

## Oracle Database

The following alerts are available for Oracle Database. Default settings for warning and alarm thresholds, duration and whether the alert is enabled (true/false) are shown.

Alert	Warning Level	Alarm Level	Duration	Enabled
<b>OraDatabaseConnectionLoss</b> Executes a single alert if the SQL database connection state is false. Index Type: Per Database Metric: Connected	NaN	NaN	30	FALSE
<b>OraDatabaseQueryError</b> Executes a single alert if the if the last query state is false (an error). Index Type: Per Database Metric: Last Query Status	NaN	NaN	30	FALSE
<b>OraDatabaseResponseTimeHigh</b> Executes a single warning and a single alarm if the time (in milliseconds) to execute a SQL query exceeds the specified threshold. Index Type: Per Database Metric: ResponseTimeMilliSec	200	220	30	FALSE
<b>OraDatabaseSpaceUsedHigh</b> Executes a single warning and a single alarm if the percent utilization of the space allocated to the database exceeds the specified threshold. Index Type: Per Database Metric: PercentUsedSpace	80	90	30	FALSE
<b>OraDatabaseTablespaceUsedHigh</b> Executes a single warning and a single alarm if the percent utilization of the database used by the tablespace exceeds the specified threshold. Index Type: Per Table Space Metric: USED_PERCENT	80	90	30	FALSE
<b>OraInstanceAvgQueryTimeHigh</b> Executes a single warning and a single alarm if the average time (in milliseconds) to perform a query exceeds the specified threshold. Index Type: Per Instance Metric: AVGQUERYTIME	300	400	30	FALSE
<b>OraInstanceCommitRateHigh</b> Executes a single warning and a single alarm if the number of commits per second exceeds the specified threshold. Index Type: Per Instance Metric: RateCOMMITTS	250	300	30	FALSE

<p><b>OraInstanceNumCurrentLoginsHigh</b> Executes a single warning and a single alarm if the number of database clients exceeds the specified threshold. Index Type: Per Instance Metric: CURRENT_LOGINS</p>	12	15	30	FALSE
<p><b>OraInstanceDataDictHitRatioLow</b> Executes a single warning and a single alarm if the data dictionary hit ratio goes below the specified threshold. Index Type: Per Instance Metric: DD_HIT_RATIO</p>	95	90	30	FALSE
<p><b>OraInstanceDiskReadRateHigh</b> Executes a single warning and a single alarm if the number of physical disk reads per second exceeds the specified threshold. Index Type: Per Instance Metric: RatePHYSICAL_READS</p>	250	300	30	FALSE
<p><b>OraInstanceDiskWriteRateHigh</b> Executes a single warning and a single alarm if the number of physical disk writes per second exceeds the specified threshold. Index Type: Per Instance Metric: RatePHYSICAL_WRITES</p>	250	300	30	FALSE
<p><b>OraInstanceLatchHitRatioLow</b> Executes a single warning and a single alarm if the latch hit ratio goes below the specified threshold. Index Type: Per Instance Metric: LatchHitPerCent</p>	95	90	30	FALSE
<p><b>OraInstanceMaxQueryTimeHigh</b> Executes a single warning and a single alarm if the query time (in milliseconds) exceeds the specified threshold. Index Type: Per Instance Metric: MAXQUERYTIME</p>	10000	15000	30	FALSE
<p><b>OraInstanceNumActiveSessionsHigh</b> Executes a single warning and a single alarm if the number of active sessions for the instance exceeds the specified threshold. Index Type: Per Instance Metric: ACTIVE_SESSIONS</p>	12	15	30	FALSE
<p><b>OraInstanceNumCurrentLoginsHigh</b> Executes a single warning and a single alarm if the number of current logins for the instance exceeds the specified threshold. Index Type: Per Instance Metric: CURRENT_LOGINS</p>	12	15	30	FALSE

<b>OraInstanceRollbackRateHigh</b> Executes a single warning and a single alarm if the number of rollbacks per second exceeds the specified threshold. Index Type: Per Instance Metric: RateROLLBACKS	5	10	30	FALSE
<b>OraInstanceSqlHitRatioLow</b> Executes a single warning and a single alarm if the SQL hit ratio goes below the specified threshold. Index Type: Per Instance Metric: SQL_HIT_RATIO	95	90	30	FALSE
<b>OraInstanceState</b> Executes a single warning and a single alarm if the database is not in an ACTIVE or OPEN state for queries. Index Type: Per Instance Metric: AlertStatus	NaN	NaN	30	FALSE

## Oracle WebLogic

The following alerts are available for Oracle WebLogic. Default settings for warning and alarm thresholds, duration and whether the alert is enabled (true/false) are shown.

Alert Name	WARN. LEVEL	ALARM LEVEL	DURATION	ENABLED
<b>WlsAppNewSessionsRateLow</b> The rate per second of newly opened sessions is below the specified threshold. <b>Index Type(s):</b> PerApplication	10	1	30	FALSE
<b>WlsAppOpenSessionsHigh</b> The maximum total number of open sessions for that application has been reached. <b>Index Type(s):</b> PerApplication	7	10	30	FALSE
<b>WlsClusterServersPercentNotRunningHigh</b> The percentage of cluster not running is high. <b>Index Type(s):</b> PerCluster	33	50	30	FALSE
<b>WlsHoggingThreadsHigh</b> The maximum number of hogging threads for that server has been reached. <b>Index Type(s):</b> PerServer	15	20	30	FALSE
<b>WlsJDBCConnectionsWaitingHigh</b> Triggered when the number of threads waiting for a JDBC connection exceeds the threshold. <b>Index Type(s):</b> PerConnection, PerLocation, PerModule, PerName	1	10	0	FALSE

<b>WlsJmsBytesCurrentHigh</b> The current number of bytes stored on this JMS server has reached its maximum. <b>Index Type(s):</b> PerServer	85	95	30	FALSE
<b>WlsJmsBytesPendingHigh</b> The current number of bytes pending (unacknowledged or uncommitted) stored on this JMS Server has reached its maximum. <b>Index Type(s):</b> PerServer	85	95	30	FALSE
<b>WlsJmsConnectionsCurrentHigh</b> The current number of connections to this JMS WebLogic Server has reached its maximum. <b>Index Type(s):</b> PerServer	85	95	30	FALSE
<b>WlsJmsDestinationBytesCurrentHigh</b> The current number of bytes stored in the destination, not including the pending bytes, has reached its maximum. <b>Index Type(s):</b> PerServer	85	95	30	FALSE
<b>WlsJmsDestinationBytesPendingHigh</b> The number of pending bytes stored in the destination has reached its maximum. <b>Index Type(s):</b> PerServer	85	95	30	FALSE
<b>WlsJmsDestinationConsumersCurrentLow</b> The number of pending bytes stored in the destination has reached its minimum. <b>Index Type(s):</b> PerServer	15	5	30	FALSE
<b>WlsJmsDestinationMessagesCurrentHigh</b> The current number of messages in the destination has reached its maximum. <b>Index Type(s):</b> PerServer	85	95	30	FALSE
<b>WlsJmsDestinationMessagesPendingHigh</b> The number of pending messages in the destination has reached its maximum. <b>Index Type(s):</b> PerServer	85	95	30	FALSE

<p><b>WlsJmsDestinationsCurrentLow</b></p> <p>The current number of destinations on this JMS Server has reached its minimum.</p> <p><b>Index Type(s):</b> PerServer</p> <p><b>Note:</b> To enable this alert, you must uncomment the following options under the <b>Collect all other metrics</b> section in the <b>sample.properties</b> file:</p> <pre># Collect all other metrics; all or none  collector.sl.rtvview.cache.config=wls_workmg r_cache.rtv collector.sl.rtvview.cache.config=wls_auxila ry_cache.rtv collector.sl.rtvview.cache.config=wls_jmsser ver_cache.rtv collector.sl.rtvview.cache.config=wls_jmsbri dge_cache.rtv collector.sl.rtvview.cache.config=wls_jmspst ore_cache.rtv</pre>	85	95	30	FALSE
<p><b>WlsJmsMessagesPendingHigh</b></p> <p>The current number of messages pending (unacknowledged or uncommitted) stored on this JMS Server has reached its maximum.</p> <p><b>Index Type(s):</b> PerServer</p>	85	95	30	FALSE
<p><b>WlsJmsServerHealthNotOK</b></p> <p>The health state of this JMS Server is not OK.</p> <p><b>Index Type(s):</b> PerServer</p>	NaN	NaN	30	FALSE
<p><b>WlsLockedUserCurrentHigh</b></p> <p>The maximum number of current locked users for that server has been reached.</p> <p><b>Index Type(s):</b> PerServer</p> <p><b>Note:</b> To enable this alert, you must uncomment the following options under the <b>Collect all other metrics</b> section in the <b>sample.properties</b> file:</p> <pre># Collect all other metrics; all or none  collector.sl.rtvview.cache.config=wls_workmg r_cache.rtv collector.sl.rtvview.cache.config=wls_auxila ry_cache.rtv collector.sl.rtvview.cache.config=wls_jmsser ver_cache.rtv collector.sl.rtvview.cache.config=wls_jmsbri dge_cache.rtv collector.sl.rtvview.cache.config=wls_jmspst ore_cache.rtv</pre>	85	95	30	FALSE
<p><b>WlsOpenSocketsHigh</b></p> <p>The maximum number of open sockets for that server has been reached.</p> <p><b>Index Type(s):</b> PerServer</p>	85	95	30	FALSE
<p><b>WlsPendingRequestCurrentHigh</b></p> <p>The maximum number of current requests for that server has been reached.</p> <p><b>Index Type(s):</b> PerServer</p>	85	95	30	FALSE

<b>WlsQueueLengthHigh</b> The number of pending requests in the priority queue has reached its maximum. This is the total of internal system requests and user requests. <b>Index Type(s):</b> PerServer	85	95	30	FALSE
<b>WlsServerCpuHigh</b> The server CPU has reached its maximum. <b>Index Type(s):</b> PerServer	85	95	30	FALSE
<b>WlsServerHealthNotOK</b> The server health is not OK. <b>Index Type(s):</b> PerServer	NaN	NaN	30	FALSE
<b>WlsServerHostCpuHigh</b> The CPU percentage of the host server has reached its maximum. <b>Index Type(s):</b> PerServer	85	95	30	FALSE
<b>WlsServerMemoryUsageHigh</b> The maximum used memory established for the server has been reached. <b>Index Type(s):</b> PerServer	85	95	30	FALSE
<b>WlsServerNewSessionsLow</b> The number of new sessions created is below the threshold. <b>Index Type(s):</b> PerServer	15	5	30	FALSE
<b>WlsServerOpenSessionsHigh</b> The maximum number of open sessions for that server has been reached. <b>Index Type(s):</b> PerServer	85	95	30	FALSE
<b>WlsServerPendingUserRequestsHigh</b> The maximum number of pending user requests has been reached. <b>Index Type(s):</b> PerServer	85	95	30	FALSE
<b>WlsServerReloadsHigh</b> The maximum number of reloads for that server has been reached. <b>Index Type(s):</b> PerServer	85	95	30	FALSE
<b>WlsServerStaleData</b> The server has stale data. <b>Index Type(s):</b> PerServer	NaN	NaN	30	FALSE
<b>WlsServerStateNotRunning</b> The state of the server is different from "Running." <b>Index Type(s):</b> PerServer	NaN	NaN	30	FALSE

<b>WlsThreadsTotalHigh</b> The total number of threads for that server has been reached. <b>Index Type(s):</b> PerServer	50	95	30	FALSE
<b>WlsTransactionRolledBackTotalHigh</b> The total number of transactions rolled back has been reached. <b>Index Type(s):</b> PerServer	85	95	30	FALSE

## RTView Host Agent

The following alerts are available for RTView Host Agent. Default settings for warning and alarm thresholds, duration and whether the alert is enabled (true/false) are shown.

Alert	Warning Level	Alarm Level	Duration	Enabled
<b>HostCpuLoadAvg1High</b> Executes a single warning alert and a single alarm alert if the average CPU load per minute exceeds the specified threshold. Index Type(s): PerHost Metric: loadAvg1	50	75	30	FALSE
<b>HostCpuLoadAvg5High</b> Executes a single warning alert and a single alarm alert if the average CPU load per 5 minutes exceeds the specified threshold. Index Type(s): PerHost Metric: loadAvg5	50	75	30	FALSE
<b>HostCpuLoadAvg15High</b> Executes a single warning alert and a single alarm alert if the average CPU load per 15 minutes exceeds the specified threshold. Index Type(s): PerHost Metric: loadAvg15	50	75	30	FALSE
<b>HostCpuPercentHigh</b> Executes a single warning alert and a single alarm alert if the percent CPU load exceeds the specified threshold. Index Type(s): PerHost Metric: hostCpuPercent	50	75	30	FALSE

<b>HostMemoryUsedHigh</b> Executes a single warning alert and a single alarm alert if the percent of physical memory used exceeds the specified threshold. Index Type(s): PerHost Metric: MemUsedPerCent	75	90	5	FALSE
<b>HostNetworkRxRateHigh</b> Executes a single warning alert and a single alarm alert if the inbound network data rate, in kilobytes per second, exceeds the specified threshold. Index Type(s): PerHost Metric: RateRxKBytes	50	75	30	FALSE
<b>HostNetworkTxRateHigh</b> Executes a single warning alert and a single alarm alert if the outbound network transmission rate, in kilobytes per second, exceeds the specified threshold. Index Type(s): PerHost Metric: RateTxKBytes	50	75	30	FALSE
<b>HostProcessCountLow</b> Executes a single warning alert and a single alarm alert if the process count exceeds the specified threshold. Index Type(s): PerHost Metric: Count	80	90	30	FALSE
<b>HostStaleData</b> Executes a single alarm alert and sets the Expired flag to <b>true</b> if data is not received from the given host within the specified expiration time interval. Index Type(s): PerHost Metric: Expired	NaN	NaN	30	FALSE
<b>HostStorageUsedHigh</b> Executes a single warning alert and a single alarm alert if the percent of space used on the storage medium exceeds the specified threshold. Index Type(s): PerStorage Metric: percentused	80	90	5	FALSE
<b>HostSwapUsedHigh</b> Executes a single warning alert and a single alarm alert if the percent of used swap space exceeds the specified threshold. Index Type(s): PerHost Metric: swapUsedPerCent	75	90	30	FALSE
<b>HostVirtualMemoryUsedHigh</b> Executes a single warning alert and a single alarm alert if the percent of used virtual memory exceeds the specified threshold. Index Type(s): PerHost Metric: VMemUsedPerCent	75	90	30	FALSE

---

## RTView Manager and RTView Rules

If RTView Manager and RTView Rules are installed on your system you might see the following alert types for RTView Servers (Data Servers, Display Servers and Historian Servers):

### RTView Server Manager Alert Types

<b>JvmCpuPercentHigh</b>	Executes a single warning alert and a single alarm alert if the percent of JVM CPU used exceeds the specified threshold. Index Type: Per JVM Metric: CpuPercent
<b>JvmGcDutyCycleHigh</b>	Executes a single warning alert and a single alarm alert if the garbage collector duty cycle exceeds the specified threshold. Index Type: Per GC Source Metric: DutyCycle
<b>JvmMemoryUsedAfterGCHigh</b>	Executes a single warning alert and a single alarm alert if the percent of memory used after garbage collection exceeds the specified threshold. Index Type: Per GC Source Metric: PctMemoryUsedAfterGC
<b>JvmMemoryUsedHigh</b>	Executes a single warning alert and a single alarm alert if the percent of memory used exceeds the specified threshold. Index Type(s): Per JVM Metric: MemoryUsedPercent
<b>JvmNotConnected</b>	Executes a single alert if the JVM is disconnected, indicating that it might have crashed. Index Type(s): Per JVM Metric: Connected
<b>JvmStaleData</b>	Executes a single alert if the data update wait time exceeds the specified duration threshold. Index Type(s): Per JVM Metric: Expired
<b>JvmThreadCountHigh</b>	Executes a single warning alert and a single alarm alert if the number of threads exceeds the specified threshold. Index Type(s): Per JVM Metric: ThreadCount
<b>TomcatAccessRateHigh</b>	Executes a single warning alert and a single alarm alert if the number of accesses per second exceeds the specified threshold. Index Type(s): Per Server Metric: RateaccessCount
<b>TomcatActiveSessionsHigh</b>	Executes a single warning alert and a single alarm alert if the number of active sessions exceeds the specified threshold. Index Type(s): Per Server Metric: activeSessions

<b>TomcatAppAccessRateHigh</b>	Executes a single warning alert and a single alarm alert if the number of accesses per second exceeds the specified threshold. Index Type(s): Per Application Metric: RateaccessCount
<b>TomcatAppActiveSessionsHigh</b>	Executes a single warning alert and a single alarm alert if the number of active sessions exceeds the specified threshold. Index Type(s): Per Application Metric: activeSessions

### RTView Rules Alert Types

<b>RtvEmServiceAlert</b>	This discrete alert is generated when a Service has one or more alerts on any associated CIs.
<b>RtvEmServiceAlertImpactHigh</b>	This limits alert is generated when a Service has an Alert Impact value that exceeds the specified threshold on any associated CI.

---

## Solace

The following alerts are available with both the solution package and standalone versions for Solace. Default settings for warning and alarm thresholds, duration and whether the alert is enabled (true/false) are shown.

Alert	Warning Level	Alarm Level	Duration	Enabled
<b>SolBridgeInboundByteRateHigh</b> The number of inbound bytes per second across the bridge has reached its maximum. Index Type: PerBridge	8000000	10000000	30	FALSE
<b>SolBridgeInboundMsgRateHigh</b> The number of inbound messages per second across the bridge as a whole has reached its maximum. Index Type: PerBridge	40000	50000	30	FALSE
<b>SolBridgeOutboundByteRateHigh</b> The number of outbound bytes per second across the bridge has reached its maximum. Index Type: PerBridge	8000000	10000000	30	FALSE
<b>SolBridgeOutboundMsgRateHigh</b> The number of outbound messages per second across the bridge has reached its maximum. Index Type: PerBridge	40000	50000	30	FALSE

<p><b>SolBrokerNoQueueFound</b>  This is an Event Alert. Event Alerts do not have duration or threshold settings.  A single alarm alert when there are discarded queues in the broker. (Delta of discard-queue-not-found is non-zero).  <b>Note:</b> This alert cannot be executed for Cloud Brokers. This request XML is a system level request which means that Cloud login credentials do not have permission to execute it.  Index Type: PerBroker</p>					FALSE
<p><b>SolBrokerNoSubscriptionMatch</b>  This is an Event Alert. Event Alerts do not have duration or threshold settings.  A single alarm alert when there are no current subscription matches (Delta of no-subscription-match is non-zero).  <b>Note:</b> This alert cannot be executed for Cloud Brokers. This request XML is a system level request which means that Cloud login credentials do not have permission to execute it.  Index Type: PerBroker</p>					FALSE
<p><b>SolBrokerNoValidDestination</b> This is an Event Alert. Event Alerts do not have duration or threshold settings.  A single alarm alert when invalid destinations exist in the broker. (Delta of discard-nodest is non-zero).  <b>Note:</b> This alert cannot be executed for Cloud Brokers. This request XML is a system level request which means that Cloud login credentials do not have permission to execute it.  Index Type: PerBroker</p>					FALSE
<p><b>SolClientInboundByteRateHigh</b>  The number of outbound bytes per second for the client has reached its maximum.  Index Type: PerClient</p>	8000000	10000000	30		FALSE
<p><b>SolClientInboundMsgRateHigh</b>  The number of outbound messages per second for the client as a whole has reached its maximum.  Index Type: PerClient</p>	40000	50000	30		FALSE
<p><b>SolClientOutboundByteRateHigh</b>  The number of outbound bytes per second for the client has reached its maximum.  Index Type: PerClient</p>	8000000	10000000	30		FALSE
<p><b>SolClientOutboundMsgRateHigh</b>  The number of outbound messages per second for the client as a whole has reached its maximum.  Index Type: PerClient</p>	40000	50000	30		FALSE
<p><b>SolClientSlowSubscriber</b>  One or more clients are consuming messages too slowly; endpoints may drop messages!  Index Type: PerClient</p>	1	NaN	30		FALSE

<b>SolCspfNeighborDown</b> State is not "OK" for one or more CSPF neighbors. Index Type: PerNeighbor	1	NaN	30	FALSE
<b>SolEndpointNoBridgeClient</b> This is an Event Alert. Event Alerts do not have duration or threshold settings. A single alarm alert when there are no binds for the Solace Endpoint exist (bind-count is zero). Index Type: PerEndpoint				FALSE
<b>SolEndpointNoBridgeTopic</b> This is an Event Alert. Event Alerts do not have duration or threshold settings. A single alarm alert when there are no topics subscribed to the Queue (topic-subscription-count is zero). Index Type: PerEndpoint				FALSE
<b>SolEndpointPendingMsgsHigh</b> The number of pending messages on a queue has reached its maximum. Index Type: PerEndpoint	8000	10000	30	FALSE
<b>SolEndpointSpoolUsageHigh</b> The endpoint is consuming too much message router memory for storing spooled messages. (Threshold units are megabytes.) Index Type: PerEndpoint	40	50	30	FALSE
<b>SolEventModuleBrokerAlert</b> This is an Event Alert. Event Alerts do not have duration or threshold settings. If the Solace Event Module is properly configured and running and this alert is enabled, all Syslog Events that are selected as alerts from the Message Brokers that were enabled for being monitored with Syslog will trigger this type of alert from the SYSTEM scope. Alerts of this type refer to Syslog events that can be clearable and non-clearable of SYSTEM scope. Therefore this alert can be clearable and non-clearable, depending on the event that triggered its execution.				FALSE
<b>SolEventModuleClientAlert</b> This is an Event Alert. Event Alerts do not have duration or threshold settings. If the Solace Event Module is properly configured and running and this alert is enabled, all Syslog Events that are selected as alerts from the Message Brokers that were enabled for being monitored with Syslog will trigger this type of alert from the CLIENT scope. Alerts of this type refer to Syslog events that can be clearable and non-clearable of CLIENT scope. Therefore this alert can be clearable and non-clearable, depending on the event that triggered its execution.				FALSE
<b>SolEventModuleVpnAlert</b> This is an Event Alert. Event Alerts do not have duration or threshold settings. If the Solace Event Module is properly configured and running and this alert is enabled, all Syslog Events that are selected as alerts from the Message Brokers that were enabled for being monitored with Syslog will trigger this type of alert from the VPN scope. Alerts of this type refer to Syslog events that can be clearable and non-clearable of VPN scope. Therefore this alert can be clearable and non-clearable, depending on the event that triggered its execution.				FALSE
<b>SolGuaranteedMsgingHbaLinkDown</b> For Guaranteed Messaging only, the Operational State for each HBA Fibre-Channel should be Online (e.g., not Linkdown). Index Type: PerHbaLink	0	NaN	30	FALSE

<b>SolGuaranteedMsgingMatePortDown</b> For Guaranteed Messaging only, the Mate Link Ports for ADB should have status OK. Index Type: PerADB	0	NaN	30	FALSE
<b>SolGuaranteedMsgingNoMsgSpoolAdActive</b> For Guaranteed Messaging only with Redundancy, at least one message router in an HA pair should show "AD-Active." Index Type: PerPair	0	NaN	30	FALSE
<b>SolMsgRouterActiveDiskUtilHigh</b> The utilization of the active disk partition for the message router is excessive. Index Type: PerAppliance	70	85	30	FALSE
<b>SolMsgRouterByteEgressUtilHigh</b> The egress rate (bytes/sec) utilization (current egress rate divided by max allowed) for the message router is excessive. Index Type: PerAppliance	70	85	30	FALSE
<b>SolMsgRouterByteIngressUtilHigh</b> The ingress rate (bytes/sec) utilization (current ingress rate divided by max allowed) for the message router is excessive. Index Type: PerAppliance	70	85	30	FALSE
<b>SolMsgRouterConnectionUtilHigh</b> The connection utilization for the message router (current number of connections divided by max allowed) is excessive. Index Type: PerAppliance	70	85	30	FALSE
<b>SolMsgRouterCpuTemperatureHigh</b> CPU temperature margin is above threshold. Index Type: PerApplianceSensor	-30	-15	30	FALSE
<b>SolMsgRouterCspfNeighborDown</b> Link-detect = no for CSPF neighbor. Index Type: PerAppliance	1	NaN	30	FALSE
<b>SolMsgRouterDelvrUnAckMsgUtilHigh</b> The delivered unacked messages as a percentage of all messages delivered for the application is excessive. Index Type: PerAppliance	70	85	30	FALSE
<b>SolMsgRouterFailoverDetected</b> The backup message router in a HA pair has assumed control. Index Type: PerAppliance	1	NaN	30	FALSE
<b>SolMsgRouterFanSensorCheckFailed</b> The speed measured for one or more fans is below threshold. Index Type: PerApplianceSensor	5000	2657	30	FALSE

<b>SolMsgRouterInboundByteRateHigh</b> The number of inbound bytes per second for the message router has reached its max threshold. Index Type: PerAppliance	400000	500000	30	FALSE
<b>SolMsgRouterInboundMsgRateHigh</b> The number of inbound messages per second for the message router has reached its max threshold. Index Type: PerAppliance	400000	500000	30	FALSE
<b>SolMsgRouterIngressFlowUtilHigh</b> The ingress flow utilization (current flows divided by max allowed) for the message router is excessive. Index Type: PerAppliance	70	85	30	FALSE
<b>SolMsgRouterInterfaceDown</b> Link-detect = no for one or more enabled network interfaces. Index Type: PerSolInterface	NaN	NaN	30	FALSE
<b>SolMsgRouterMsgCountUtilHigh</b> The message count utilization for the message router is excessive. Index Type: PerAppliance	70	85	30	FALSE
<b>SolMsgRouterMsgEgressUtilHigh</b> The message egress rate utilization (current message egress rate divided by max allowed) for the message router is excessive. Index Type: PerAppliance	70	85	30	FALSE
<b>SolMsgRouterMsgIngressUtilHigh</b> The message ingress rate utilization (current message ingress rate divided by max allowed) for the message router is excessive. Index Type: PerAppliance	70	85	30	FALSE
<b>SolMsgRouterNABUsageHigh</b> Network Acceleration Blade memory usage is excessive. Index Type: PerNAB	60	80	30	FALSE
<b>SolMsgRouterNotConnected</b> The message router is not ready for collecting performance monitoring data. Index Type: PerAppliance	NaN	NaN	30	FALSE
<b>SolMsgRouterOutboundByteRateHigh</b> The number of outbound bytes per second for the message router has reached its max threshold. Index Type: PerAppliance	400000	500000	30	FALSE
<b>SolMsgRouterOutboundMsgRateHigh</b> The number of outbound messages per second for the message router has reached its max threshold. Index Type: PerAppliance	400000	500000	30	FALSE
<b>SolMsgRouterPendingMsgsHigh</b> The total number of pending messages for this message router has reached its maximum. Index Type: PerAppliance	400000	500000	30	FALSE

<b>SolMsgRouterPowerSupplyFailed</b> A power supply has failed. Index Type: PerAppliance	0	NaN	30	FALSE
<b>SolMsgRouterSpoolUtilization</b> The amount of spool space used for messages is excessive. Index Type: PerAppliance	70	85	30	FALSE
<b>SolMsgRouterStandbyDiskUtilHigh</b> The utilization of the standby disk partition for the message router is excessive. Index Type: PerAppliance	70	85	30	FALSE
<b>SolMsgRouterSubscriptionUtilHigh</b> The subscription utilization (current number of subscriptions divided by max allowed) for the message router is excessive. Index Type: PerAppliance	70	85	30	FALSE
<b>SolMsgRouterSwapUsedHigh</b> The amount of swap space used by the message router operating system is excessive. Index Type: PerAppliance	70	85	30	FALSE
<b>SolMsgRouterSyslogAlert</b> This alert executes when a Solace Syslog Warning or Critical message is received. To get Syslog event alerts (in RTView Enterprise or the standalone Monitor), go to the Alert Administration display and enable the <b>SolMsgRouterSyslog</b> alert.	-	-	-	-
<b>SolMsgRouterTemperatureSensorCheckFailed</b> A chassis temperature measurement is above threshold. Index Type: PerAppliance	40	45	30	FALSE
<b>SolMsgRouterTranSessionCntUtilHigh</b> The transacted session count utilization for the message router is excessive. Index Type: PerAppliance	70	85	30	FALSE
<b>SolMsgRouterTranSessionResUtilHigh</b> The transacted session resource utilization for the message router is excessive. Index Type: PerAppliance	70	85	30	FALSE
<b>SolMsgRouterVoltageSensorCheckFailed</b> A power supply voltage is high or low. Index Type: PerApplianceSesor	NaN	NaN	30	FALSE
<b>SolSparseMessageSpoolFile</b> This is a Limits Alert that issues a Warning alert and is enabled by default. <b>Important:</b> Do not modify thresholds for this alert as they were set up by Solace development team) This alert is defined to determine when there is a Sparse Message Spool File Condition. When disk space usage is several multiples of persistent store usage, then there is likely a large number of message spool files residing on the disk where each file contains few messages. This is referred to as a sparse message spool file condition, and requires urgent attention to mitigate and avoid the disk reaching capacity. For further information, refer to Solace documentation.				TRUE

<b>SolVpnConnectionCountHigh</b> The number of connections to the server has reached its maximum. Index Type: PerVPN	60	80	30	FALSE
<b>SolVpnInboundByteRateHigh</b> The number of inbound bytes per second for the vpn has reached its maximum. Index Type: PerVPN	8000000	10000000	30	FALSE
<b>SolVpnInboundDiscardRateHigh</b> The number of discarded inbound messages per second for the server is excessive. Index Type: PerVPN	1	5	30	FALSE
<b>SolVpnInboundMsgRateHigh</b> The number of inbound messages per second for the vpn as a whole has reached its maximum. Index Type: PerVPN	40000	50000	30	FALSE
<b>SolVpnOutboundByteRateHigh</b> The number of outbound bytes per second for the VPN has reached its maximum. Index Type: PerVPN	8000000	10000000	30	FALSE
<b>SolVpnOutboundDiscardRateHigh</b> The number of discarded outbound messages per second for the server is excessive. Index Type: PerVPN	1	5	30	FALSE
<b>SolVpnOutboundMsgRateHigh</b> The number of outbound messages per second for the server as a whole has reached its maximum. Index Type: PerVPN	40000	50000	30	FALSE
<b>SolVpnPendingMsgsHigh</b> The total number of pending messages for this destination has reached its maximum. Index Type: PerVPN	8000000	10000000	30	FALSE
<b>SolVpnSubscriptionCountHigh</b> The number of endpoints in this VPN has reached its maximum. Index Type: PerVPN	8000	10000	30	FALSE

## TIBCO ActiveMatrix BusinessWorks

The following alerts are available with both the solution package and standalone versions for TIBCO® ActiveMatrix BusinessWorks™. Default settings for warning and alarm thresholds, duration and whether the alert is enabled (true/false) are shown.

Alert	Warning Level	Alarm Level	Duration	Enabled
<b>BW6AppErrorState</b> BW6 application status is not Running or Stopped (status is Impaired, AppError or StartFailed) Index Type: PerApp Metric: State	NaN	NaN	30	FALSE
<b>Bw6AppExpired</b> BW6 application expired due to application inactivity. Index Type: PerApp Metric: Stopped	NaN	NaN	30	FALSE
<b>Bw6AppNodeCpuUsedHigh</b> BW6 AppNode CPU usage exceeded limit. CPU Usage is the CPU time in use by all processes expressed as a percentage of the total CPU time available. Index Type: PerAppNode Metric: CPU Usage%	50	80	30	FALSE
<b>Bw6AppNodeMemUsedHigh</b> BW6 AppNode memory usage exceeded limit. Memory usage is the percentage of total JVM memory currently consumed by this appnode. Index Type: PerAppNode Metric: Memory Usage%	50	80	30	FALSE
<b>Bw6AppNodeStopped</b> BW6 AppNode stopped purposefully (for example, an administrator stopped the AppNode process). Index Type: PerAppNode Metric: State	NaN	NaN	10	FALSE
<b>Bw6AppNodeUnreachable</b> BW6 AppNode stopped abnormally (for example, the AppNode process crashed). Index Type: PerAppNode Metric: State	NaN	NaN	10	FALSE
<b>Bw6AppProcessCreatedRateHigh</b> BW6 Process created rate for application exceeded limit. Index Type: PerApp Metric: App Created Rate	50	80	30	FALSE
<b>Bw6AppProcessElapsedTimeHigh</b> BW6 Process delta elapsed time rate of increase for application exceeded limit. Index Type: PerApp Metric: App Elapsed Rate	200	400	30	FALSE

<b>Bw6AppProcessExecutionTimeHigh</b> BW6 Process delta execution time rate of increase for application exceeded limit. Index Type: PerApp Metric: App Execution Rate	200	400	30	FALSE
<b>Bw6AppProcessFailedRateHigh</b> BW6 Process failed rate for application exceeded limit. Index Type: PerApp Metric: App Failed Rate	50	80	30	FALSE
<b>Bw6AppStopped</b> BW6 application stopped. Index Type: PerApp Metric: Stopped	NaN	NaN	30	FALSE
<b>Bw6ProcessActivityErrorRateHigh</b> BW6 Process error rate exceeded limit. Index Type: PerProcess Metric: Process Failed Rate	50	80	30	FALSE
<b>Bw6ProcessCreatedRateHigh</b> BW6 Process error rate exceeded limit. Index Type: PerProcess Metric: Process Failed Rate	50	80	30	FALSE
<b>Bw6ProcessElapsedTimeHigh</b> BW6 Process delta elapsed time rate of increase exceeded limit. Index Type: PerProcess Metric: Delta Exec Rate	200	400	30	FALSE
<b>Bw6ProcessExecutionTimeHigh</b> BW6 Process delta execution time rate of increase exceeded limit. Index Type: PerProcess Metric: Delta Time Rate	200	400	30	FALSE
<b>Bw6ProcessFailedRateHigh</b> BW6 Process suspended rate exceeded limit. Index Type: PerProcess Metric: Suspended Rate	50	80	30	FALSE
<b>Bw6ProcessHung</b> The delta elapsed time is greater than zero but the delta execution time is zero. Index Type: PerProcess Metric: Hung/Not Hung	NaN	NaN	10	FALSE
<b>Bw6ProcessSuspendRateHigh</b> BW6 Process failed rate exceeded limit. Index Type: PerProcess Metric: Failed Rate	50	80	30	FALSE

<b>BwActivityErrorRateHigh</b> BW5 Activity error rate exceeded limit. The rate is calculated by taking the delta of total error returns in this update period and dividing by the length of the period. Index Type: PerActivity Metric: RateErrorCount	50	80	30	FALSE
<b>BwActivityExecutionTimeHigh</b> BW5 Activity execution time rate of increase exceeded limit The rate is calculated by taking the delta of total execution time in this update period and dividing by the length of the period. Index Type: PerActivity Metric: RateExecutionTime	200	400	30	FALSE
<b>BwEngineCpuUsedHigh</b> BW Engine CPU usage (% of total) exceeded limit. CPU Usage is the CPU time used by the BW engine expressed as a percentage of the total CPU time available. Index Type: PerEngine Metric: CPU Usage%	50	80	30	FALSE
<b>BwEngineMemUsedHigh</b> BW Engine memory usage (% of total) exceeded limit. Memory usage is the percentage of total JVM memory currently consumed by this engine. Index Type: PerEngine Metric: PercentUsed	50	80	30	FALSE
<b>BwEngineStopped</b> BW Engine has stopped running. Index Type: PerEngine Metric: Stopped	NaN	NaN	30	FALSE
<b>BwEngineUnreachable</b> BW engine stopped abnormally. Index Type: PerEngine Metric: State	NaN	NaN	30	FALSE
<b>BwProcessAbortRateHigh</b> BW Process aborted rate exceeded limit. The rate is calculated by taking the delta of total aborts in this update period and dividing by the length of the period. Index Type: PerProcess Metric: RateAborted	50	80	30	FALSE
<b>BwProcessAvgElapsedTimeHigh</b> BW Process Average Elapsed Time exceeded limit. Value is calculated by dividing the delta elapsed time for the interval by the delta completed, or the number of process instances that completed in the interval. Index Type: PerProcess Metric: Process Avg Elapsed Time	100	200	30	FALSE
<b>BwProcessAvgExecutionTimeHigh</b> BW Process average execution time exceeded limit. Index Type: PerProcess Metric: AverageExecution	0	0	0	FALSE

<p><b>BwProcessCreatedRateHigh</b>          BW Process creation rate exceeded limit. The rate is calculated by taking the number of process instances created in the interval and dividing by the length of the interval in seconds.          Index Type: PerProcess          Metric: Processes Created/sec</p>	100	200	30	FALSE
<p><b>BwProcessCreatedRateLow</b>          BW Process creation rate per second went below limit.          Index Type: PerProcess          Metric: App Created Rate</p>	0	0	0	FALSE
<p><b>BwProcessElapsedTimeHigh</b>          BW Process elapsed time rate of increase exceeded limit. The rate is calculated by taking the delta of total elapsed time in this update period and dividing by the length of the period.          Index Type: PerProcess          Metric: RateTotalElapsed</p>	50	80	30	FALSE
<p><b>BwProcessExecutionTimeHigh</b>          BW Process execution time rate of increase exceeded limit. The rate is calculated by taking the delta of total execution time in this update period and dividing by the length of the period.          Index Type: PerProcess          Metric: RateTotalExecution</p>	50	80	30	FALSE
<p><b>BwProcessHung</b>          The delta elapsed time is greater than zero but the delta execution time is zero.          Index Type: PerProcess          Metric: Hung/Not Hung</p>	NaN	NaN	10	FALSE
<p><b>BwProcessSuspendRateHigh</b>          BW Process suspended rate exceeded limit. The rate is calculated by taking the delta of total suspends in this update period and dividing by the length of the period.          Index Type: PerProcess          Metric: RateSuspended</p>	50	80	30	FALSE
<p><b>BwProcessTotalCpuPercentHigh</b>          BW Process CPU percent utilization exceeded limit. This is the percent CPU used by all process instances executing over the interval.          Index Type: PerProcess          Metric: Process Total CPU Percent</p>	50	75	30	FALSE
<p><b>BwServerCpuUsedHigh</b>          BW Server CPU usage (% of total) exceeded limit. CPU Usage is the CPU time in use by all processes expressed as a percentage of the total CPU time available.          Index Type: PerServer          Metric: CPU Usage%</p>	60	85	30	FALSE

<p><b>BwServerFreeMemLow</b>          BW Server free memory available (in megabytes) is below limit. Free memory means available physical (RAM) memory.          Index Type: PerServer          Metric: Memory Free Mbytes</p>	15	5	30	FALSE
<p><b>BwServerInactive</b>          BW Server has become inactive. The period of time specified by the substitution variable \$bwserverExpirationTime has passed since data was last received from the server.          Index Type: PerServer          Metric: Expired</p>	NaN	NaN	30	FALSE
<p><b>BwServerMemUsedHigh</b>          BW Server memory usage (% of total) exceeded limit. Memory usage is the virtual memory in use expressed as a percentage of the available virtual memory. The meaning of available virtual memory is system-dependent: on Windows it refers to pagefile space; on Unix systems it refers to swap space.          Index Type: PerServer          Metric: Virtual Memory Used%</p>	50	80	30	FALSE
<p><b>HawkAlert</b>          Display Hawk alerts throughout the Monitor. To enable Hawk Alerts to be included in alert counts and displayed throughout the Monitor, scroll down to <b>HawkAlert</b> in the <b>Active Alert Table</b> and select the <b>Alert Enabled</b> checkbox. It is possible to filter unwanted alerts from the cache data so that those alerts are not included throughout the Monitor.          To filter unwanted alerts out of the Hawk cache data, enter the following into the <b>sample.properties</b> file (located in the project directory you created). NOTE: Unwanted alerts are filtered out according to the AlertText.  <b>sl.rtvew.sub=\$hawkAlertTextFilterOut:AlertText</b>          For example, to filter out all Hawk Alerts in which the AlertText contains <b>Source</b> you would enter the following:  <b>sl.rtvew.sub=\$hawkAlertTextFilterOut:Source</b>          The default time to remove cleared Hawk Alerts from the table is <b>3600</b> seconds. To adjust this setting, edit the following in <b>sample.properties</b>:  <b>sl.rtvew.sub=\$hawkAlertTextFilterOut:3600</b>          Index Type: PerServer          Metric: Hawk</p>	NaN	NaN	-1	TRUE
<p><b>JvmCpuPercentHigh</b>          The percentage of CPU that has been reached by the JVM is above the limit.          Index Type: PerJVM          Metric: CpuPercent</p>	50	75	30	FALSE
<p><b>JvmGcDutyCycleHigh</b>          The duty cycle is out the upper limit.          Index Type: PerGC          Metric: DutyCycle</p>	50	75	30	FALSE

<b>JvmMemoryUsedHigh</b> The memory used out the upper limit Index Type: PerJVM Metric: MemoryUsedPercent	50	75	30	FALSE
<b>JvmNotConnected</b> The JVM in not connected. Index Type: PerJVM Metric: Connected	NaN	NaN	30	FALSE
<b>JvmStaleData</b> Cut in reception from that JVM. Index Type: PerJVM Metric: Expired	NaN	NaN	30	FALSE

## TIBCO ActiveSpaces

The following alerts are available for TIBCO ActiveSpaces. Default settings for warning and alarm thresholds, duration and whether the alert is enabled (true/false) are shown.

Alert Name	WARN. LEVEL	ALARM LEVEL	DURATION	ENABLED
<b>TdgKeeperCpuUsageHigh</b> The keeper CPU usage rate (msec/sec) is above the defined thresholds. <b>Index Type(s):</b> PerTdgKeeper	60	80	30	FALSE
<b>TdgKeeperExpired</b> RTView is not receiving metrics updates from this Keeper. The Expired flag of the Keeper was set to true. <b>Index Type(s):</b> PerTdgKeeper	NaN	NaN	30	FALSE
<b>TdgKeeperMemoryUseHigh</b> The keeper's usage of memory, in KB, is above the threshold. <b>Index Type(s):</b> PerTdgKeeper	1600000	2000000	30	FALSE
<b>TdgKeeperMsgsRcvdRateHigh</b> The incoming message rate, in messages per second, is higher than expected for this keeper. <b>Index Type(s):</b> PerTdgKeeper	160000	200000	30	FALSE
<b>TdgKeeperMsgsSentRateLow</b> The keeper's rate of messages sent is below the threshold. <b>Index Type(s):</b> PerTdgKeeper	15	5	30	FALSE

<b>TdgKeeperNotRunning</b> The current status for this keeper is not "RUNNING." <b>Index Type(s):</b> PerTdgKeeper	NaN	NaN	30	FALSE
<b>TdgNodeCpuUsageHigh</b> The node CPU Usage rate (msec/sec) is above threshold. <b>Index Type(s):</b> PerTdgNode	60	80	30	FALSE
<b>TdgNodeExpired</b> RTView is not receiving metrics updates from this Node. The Expired flag of the Node was set to true. <b>Index Type(s):</b> PerTdgNode	NaN	NaN	30	FALSE
<b>TdgNodeLiveDataSizeHigh</b> The node's live data size is above the threshold. <b>Index Type(s):</b> PerTdgNode	1600000	2000000	30	FALSE
<b>TdgNodeMemoryUseHigh</b> The node's usage of memory, in KB, is above the threshold. <b>Index Type(s):</b> PerTdgNode	1600000	2000000	30	FALSE
<b>TdgNodeMsgsRcvdRateHigh</b> The incoming message rate, in messages per second, is higher than expected for this node. <b>Index Type(s):</b> PerTdgNode	160000	200000	30	FALSE
<b>TdgNodeMsgsSentRateLow</b> The outgoing message rate, in messages per second, is lower than expected for this node. <b>Index Type(s):</b> PerTdgNode	15	5	30	FALSE
<b>TdgNodeNotRunning</b> The current status for this node is not "RUNNING". <b>Index Type(s):</b> PerTdgNode	NaN	NaN	30	FALSE
<b>TdgNodeOpsCompletedRateLow</b> The rate of completed operations on the node is below the threshold. <b>Index Type(s):</b> PerTdgNode	15	5	30	FALSE
<b>TdgNodeOpsFailedRateHigh</b> The rate of failed operations on the node is above the threshold. <b>Index Type(s):</b> PerTdgNode	10	20	30	FALSE
<b>TdgNodeTxnRollbackRateHigh</b> The node's rate of transactions rolled back is above the threshold. <b>Index Type(s):</b> PerTdgNode	50	100	30	FALSE
<b>TdgProxyCpuUsageHigh</b> The proxy CPU Usage rate (msec/sec) is above the defined threshold. <b>Index Type(s):</b> PerTdgProxy	60	80	30	FALSE

<b>TdgProxyExpired</b> RTView is not receiving metrics updates from this Proxy. The Expired flag of the Proxy was set to true. <b>Index Type(s):</b> PerTdgProxy	NaN	NaN	30	FALSE
<b>TdgProxyMemoryUseHigh</b> The proxy's usage of memory, in kilobytes, is above the threshold. <b>Index Type(s):</b> PerTdgProxy	1600000	2000000	30	FALSE
<b>TdgProxyMsgsRcvdRateHigh</b> The incoming message rate, in messages per second, is higher than expected for this proxy. <b>Index Type(s):</b> PerTdgProxy	160000	200000	30	FALSE
<b>TdgProxyMsgsSentRateLow</b> The outgoing message rate, in messages per second, is lower than expected for this proxy. <b>Index Type(s):</b> PerTdgProxy	15	5	30	FALSE
<b>TdgProxyNotRunning</b> The current status for this proxy is not "RUNNING." <b>Index Type(s):</b> PerTdgProxy	NaN	NaN	30	FALSE
<b>TdgProxyTxnRollbackRateHigh</b> The proxy's rate of transactions rolled back is above the threshold. <b>Index Type(s):</b> PerTdgProxy	50	100	30	FALSE
<b>TdgRealmOpsCompletedRateLow</b> The rate of completed operations on the realm is below the threshold. <b>Index Type(s):</b> PerTdgRealm	15	5	30	FALSE
<b>TdgRealmOpsFailedRateHigh</b> The rate of failed operations on the realm is above the threshold. <b>Index Type(s):</b> PerTdgRealm	10	20	30	FALSE
<b>TdgRealmServerCpuUsageHigh</b> The CPU utilization of the Realm Server, as a percentage, is above the threshold. <b>Index Type(s):</b> PerTdgRealm	60	80	30	FALSE
<b>TdgRealmServerExpired</b> RTView is not receiving metrics updates from this Realm Server. The Expired flag was set to true. <b>Index Type(s):</b> PerTdgRealm	NaN	NaN	30	FALSE
<b>TdgRealmServerMemoryUseHigh</b> The Realm Server memory usage (RSS) is above threshold. Units are kilobytes. <b>Index Type(s):</b> PerTdgRealm	160	200	30	FALSE
<b>TdgRealmTxnRollbackRateHigh</b> The node's rate of transactions rolled back is above the threshold. <b>Index Type(s):</b> PerTdgRealm	50	100	30	FALSE

## TIBCO ActiveSpaces (2.x)

The following alerts are available for TIBCO ActiveSpaces (2.x). Default settings for warning and alarm thresholds, duration and whether the alert is enabled (true/false) are shown.

Alert Name	WARN. LEVEL	ALARM LEVEL	DURATION	ENABLED
<b>TasMemberCpuHigh</b> The CPU usage is above the defined thresholds. <b>Index Type(s):</b> PerMember	80	95	30	FALSE
<b>TasMemberEntriesHigh</b> The number of objects inserted into the space is above the defined thresholds. <b>Index Type(s):</b> PerMember	8000	10000	30	FALSE
<b>TasMemberEvictsRateHigh</b> The rate at which 'evicts' are occurring is above the defined thresholds. <b>Index Type(s):</b> PerMember	80	100	30	FALSE
<b>TasMemberExpireRateHigh</b> The rate at which 'expires' are occurring is above the defined thresholds. <b>Index Type(s):</b> PerMember	80	100	30	FALSE
<b>TasMemberGetRateHigh</b> The rate at which 'gets' are occurring is above the defined thresholds. <b>Index Type(s):</b> PerMember	80	100	30	FALSE
<b>TasMemberJvmMemoryUsedHigh</b> The percent JVM memory used is above the defined thresholds. <b>Index Type(s):</b> PerMember	80	95	30	FALSE
<b>TasMemberMemoryUsedHigh</b> The percent memory used is above the defined thresholds. <b>Index Type(s):</b> PerMember	80	95	30	FALSE
<b>TasMemberPutRateHigh</b> The rate at which 'puts' are occurring is above the defined thresholds. <b>Index Type(s):</b> PerMember	80	100	30	FALSE

<b>TasMemberSeederCapacity</b> The percentage utilization (number of entries/capacity)*100 of the seeder is high for the given space. "Capacity per seeder" must be set in the space definition for this alarm to be effective. <b>Index Type(s):</b> PerMemberandSpace	80	90	30	FALSE
<b>TasMemberTakeRateHigh</b> The rate at which 'takes' are occurring is above the defined thresholds. <b>Index Type(s):</b> PerMember	80	100	30	FALSE
<b>TasMetaspaceEntriesHigh</b> The number of objects inserted into the metaspace is above the defined thresholds. <b>Index Type(s):</b> PerMetaspace	8000	100000	30	FALSE
<b>TasMetaspaceEvictsRateHigh</b> The rate at which 'evicts' are occurring is above the defined thresholds. <b>Index Type(s):</b> PerMetaspace	80	100	30	FALSE
<b>TasMetaspaceExpireRateHigh</b> The rate at which 'expires' are occurring is above the defined thresholds. <b>Index Type(s):</b> PerMetaspace	80	100	30	FALSE
<b>TasMetaspaceGetRateHigh</b> The rate at which 'gets' are occurring is above the defined thresholds. <b>Index Type(s):</b> PerMetaspace	80	100	30	FALSE
<b>TasMetaspacePutRateHigh</b> The rate at which 'puts' are occurring is above the defined thresholds. <b>Index Type(s):</b> PerMetaspace	80	100	30	FALSE
<b>TasMetaspaceTakeRateHigh</b> The rate at which 'takes' are occurring is above the defined thresholds. <b>Index Type(s):</b> PerMetaspace	80	100	30	FALSE
<b>TasQueryDurationHigh</b> The query duration (in seconds) is above the defined threshold (in seconds). <b>Index Type(s):</b> PerSpace	4	5	30	FALSE
<b>TasSpaceEntriesHigh</b> The number of objects inserted into the space is above the defined thresholds. <b>Index Type(s):</b> PerSpace	8000	100000	30	FALSE
<b>TasSpaceEvictsRateHigh</b> The rate at which 'evicts' are occurring is above the defined thresholds. <b>Index Type(s):</b> PerSpace	80	100	30	FALSE

<b>TasSpaceExpireRateHigh</b> The rate at which 'expires' are occurring is above the defined thresholds. <b>Index Type(s):</b> PerSpace	80	100	30	FALSE
<b>TasSpaceGetRateHigh</b> The rate at which 'gets' are occurring is above the defined thresholds. <b>Index Type(s):</b> PerSpace	80	100	30	FALSE
<b>TasSpacePutRateHigh</b> The rate at which 'puts' are occurring is above the defined thresholds. <b>Index Type(s):</b> PerSpace	80	100	30	FALSE
<b>TasSpaceSeederCountLow</b> Not enough seeders are available. <b>Index Type(s):</b> PerSpace	NaN	NaN	30	FALSE
<b>TasSpaceState</b> The state of the space is "not ready". <b>Index Type(s):</b> PerSpace	NaN	NaN	30	FALSE
<b>TasSpaceTakeRateHigh</b> The rate at which 'takes' are occurring is above the defined thresholds. <b>Index Type(s):</b> PerSpace	80	100	30	FALSE

## TIBCO Adapters

The following alerts are available for TIBCO Adapters. Default settings for warning and alarm thresholds, duration and whether the alert is enabled (true/false) are shown.

<b>Alert Name</b>	<b>WARN. LEVEL</b>	<b>ALARM LEVEL</b>	<b>DURATION</b>	<b>ENABLED</b>
<b>TadAdapterDeltaErrorsHigh</b> The number of errors incurred by the adapter in last measurement interval is above the defined threshold. <b>Index Type(s):</b> PerAdapter	1600	2000	30	FALSE
<b>TadAdapterExpired</b> The data from this adapter has not been updated since the last measurement interval. The data shown from this adapter is currently stale. <b>Index Type(s):</b> PerAdapter	NaN	NaN	0	FALSE

<b>TadAdapterMsgsRcvdRateHigh</b>	1600	2000	60	FALSE
The number of messages received by this adapter since the last measurement interval is above the defined threshold.				
<b>Index Type(s):</b> PerAdapter				
<b>TadAdapterMsgsSentRateHigh</b>	1	2	60	FALSE
The number of messages sent by the adapter in the last measurement interval is above the defined threshold.				
<b>Index Type(s):</b> PerAdapter				

---

## TIBCO BusinessEvents

The following alerts are available with both the solution package and standalone versions for TIBCO® BusinessEvents®.

<b>TbeBackingStoreEraseRateHigh</b>	This alert executes a single warning alert and a single alarm alert if the rate at which entries are erased from the backing store exceeds the specified threshold. The warning default threshold is <b>80</b> and the alarm default threshold is <b>95</b> .
<b>TbeBackingStoreLoadRateHigh</b>	This alert executes a single warning alert and a single alarm alert if the rate at which entries are loaded from the backing store exceeds the specified threshold. The warning default threshold is <b>80</b> and the alarm default threshold is <b>95</b> .
<b>TbeBackingStoreStoreRateHigh</b>	This alert executes a single warning alert and a single alarm alert if the rate at which entries are written to the backing store exceeds the specified threshold. The warning default threshold is <b>80</b> and the alarm default threshold is <b>95</b> .
<b>TbeClusterMalformed</b>	<p>This alert executes for any cluster where the member count is not equal to the expected cluster size. The expected cluster size is a count of the number of nodes that have the same cluster name, as discovered by reading the cluster MBean for each node in the connection property file. The MemberCount attribute is also read from the same cluster MBean, and is the number of nodes in the (sub)cluster which the current node has joined.</p> <p>The condition where these counts differ can occur if there are missing connections in the property file (for example, some nodes are unmonitored). It can also occur if, due to network or other anomalies, some nodes do not join the "main" cluster, but instead form a "sub-cluster" of one or more nodes. This condition is commonly referred to as "split-brain".</p>
<b>TbeDestinationStatusRecvdEventsRateHigh</b>	This alert executes a single warning alert and a single alarm alert if the rate at which events are received from the channel exceeds the specified threshold. The warning default threshold is <b>80</b> and the alarm default threshold is <b>95</b> .
<b>TbeNodeConceptsGetRateHigh</b>	This alert executes a single warning alert and a single alarm alert if the rate at which concepts are received from the cache exceeds the specified threshold. The warning default threshold is <b>80</b> and the alarm default threshold is <b>95</b> .

<b>TbeNodeConceptsPutRateHigh</b>	This alert executes a single warning alert and a single alarm alert if the rate at which concepts are written to the cache exceeds the specified threshold. The warning default threshold is <b>80</b> and the alarm default threshold is <b>95</b> .
<b>TbeNodeConceptsRemoveRateHigh</b>	This alert executes a single warning alert and a single alarm alert if the rate which concepts are removed from the cache exceeds the specified threshold. The warning default threshold is <b>80</b> and the alarm default threshold is <b>95</b> .
<b>TbeNodeConnectionLoss</b>	This discrete alert executes when the JMX Connection to the TIBCO BusinessEvents agent is lost (the TCP connection flag for an engine is <b>false</b> ).
<b>TbeNodeEventsGetRateHigh</b>	This alert executes a single warning alert and a single alarm alert if the rate at which events are received from the cache exceeds the specified threshold. The warning default threshold is <b>80</b> and the alarm default threshold is <b>95</b> .
<b>TbeNodeEventsPutRateHigh</b>	This alert executes a single warning alert and a single alarm alert if the rate at which events are written to the cache exceeds the specified threshold. The warning default threshold is <b>80</b> and the alarm default threshold is <b>95</b> .
<b>TbeNodeEventsRemoveRateHigh</b>	This alert executes a single warning alert and a single alarm alert if the rate which events are removed from the cache exceeds the specified threshold. The warning default threshold is <b>80</b> and the alarm default threshold is <b>95</b> .
<b>TbeObjectTableExtIdSize</b>	This alert executes a single warning alert and a single alarm alert if the number of external object IDs exceeds the specified threshold. The warning default threshold is <b>9000</b> and the alarm default threshold is <b>10000</b> .
<b>TbeObjectTableSize</b>	This alert executes a single warning alert and a single alarm alert if the number of objects maintained by the cache exceeds the specified threshold. The warning default threshold is <b>9000</b> and the alarm default threshold is <b>10000</b> .
<b>TbeRuleFiringRateHigh</b>	This alert executes a single warning alert and a single alarm alert if the rate at which rules are executing exceeds the specified threshold. The warning default threshold is <b>80</b> and the alarm default threshold is <b>95</b> .

## TIBCO Enterprise Message Service

The following alerts are available with both the solution package and standalone versions for TIBCO® Enterprise Message Service™. Default settings for warning and alarm thresholds, duration and whether the alert is enabled (true/false) are shown.

Alert Name	WARN. LEVEL	ALARM LEVEL	DURATION	ENABLED
<p><b>EmsConsumerStalled</b></p> <p>Indicates consumers are stalled or are no longer consuming messages (have not received a message within a defined threshold). The server must be running for a minimum time (5 minutes by default) before this alert is triggered. Thresholds are in seconds.</p> <p><b>Note:</b> This alert does not allow overrides.</p> <p><b>Index Type(s):</b> PerConsumer:ID/ PerServerConsumer:URL;ID</p> <p><b>Metric:</b> elapsedSinceLasAckInSec</p>	85	95	30	FALSE
<p><b>EmsConsumerStuck</b></p> <p>Indicates a consumer is stuck because there are existing messages that can be consumed (currentMsSentCount &gt; 0), but none of the messages have been consumed within the defined warning and alert thresholds (elapsedSinceLasAckInSec &gt; threshold). Alert and warning thresholds are in seconds.</p> <p><b>Index Type(s):</b> PerConsumer:ID/ PerServerConsumer:URL;ID</p> <p><b>Metric:</b> currentMsgSentCount, elapsedSinceLasAckInSec</p>	85	95	30	FALSE
<p><b>EmsQueueConsumerIdleTimeHigh</b></p> <p>The idle time of the queue consumer has reached its maximum. This alert is triggered when there is no change in the number of incoming messages for a queue within a specified period of time (in seconds).</p> <p><b>Index Type(s):</b> PerQueue;PerServerQueue</p> <p><b>Metric:</b> ConsumerIdleTime</p>	60	80	30	FALSE
<p><b>EmsQueueInboundDeltaHigh</b></p> <p>The number of new incoming messages for the EMS Queue has reached its maximum.</p> <p><b>Index Type(s):</b> PerQueue;PerServerQueue</p> <p><b>Metric:</b> DeltainboundTotalMessages</p>	60	80	30	FALSE
<p><b>EmsQueueMsgLatencyHigh</b></p> <p>The time, in seconds, needed to process all pending messages based on the current outbound message rate exceeded its threshold. This alert does not take into account queues with outbound message rate equals to zero.</p> <p><b>Index Type(s):</b> PerServerQueue:URL;name</p> <p><b>Metric:</b> messageLatency</p>	60	80	30	FALSE

<p><b>EmsQueueProviderIdleTimeHigh</b></p> <p>The queue idle time exceeded the specified threshold. A queue is idle when the number of inbound messages remains unchanged.</p> <p><b>Index Type(s):</b> PerServerQueue:URL;name</p> <p><b>Metric:</b> ProviderIdleTime</p>	60	80	30	FALSE
<p><b>EmsQueuesConsumerCountHigh</b></p> <p>The number of consumers of a queue exceeded the specified high threshold.</p> <p><b>Index Type(s):</b> PerServerQueue:URL;name/ PerQueue:name</p> <p><b>Metric:</b> consumerCount</p>	60	80	30	FALSE
<p><b>EmsQueuesConsumerCountLow</b></p> <p>The number of consumers of a queue is below the specified threshold.</p> <p><b>Index Type(s):</b> PerServerQueue:URL;name/ PerQueue:name</p> <p><b>Metric:</b> consumerCount</p>	15	5	30	FALSE
<p><b>EmsQueuesInMsgRateHigh</b></p> <p>The rate of inbound messages on the queue exceeded the specified threshold.</p> <p><b>Index Type(s):</b> PerServerQueue:URL;name/ PerQueue:name</p> <p><b>Metric:</b> inboundMessageRate</p>	60	80	30	FALSE
<p><b>EmsQueuesOutMsgRateHigh</b></p> <p>The number of outbound messages on the queue exceeded the specified threshold.</p> <p><b>Index Type(s):</b> PerServerQueue:URL;name</p> <p><b>Metric:</b> outboundMessageRate</p>	60	80	30	FALSE
<p><b>EmsQueuesPendingMsgsHigh</b></p> <p>The number of pending messages on the queue exceeded the specified threshold.</p> <p><b>Index Type(s):</b> PerServerQueue:name;PerServerQueue:URL;name</p> <p><b>Metric:</b> pendingMessageCount</p>	60	80	30	FALSE
<p><b>EmsQueuesProducerCountHigh</b></p> <p>The number of producers to a queue exceeded the specified high threshold.</p> <p><b>Index Type(s):</b> PerQueue:name/ PerServerQueue:URL;name</p> <p><b>Metric:</b> producerCount</p>	60	80	30	TRUE
<p><b>EmsQueuesProducerCountLow</b></p> <p>The number of producers to a queue is below the specified threshold.</p> <p><b>Index Type(s):</b> PerQueue:name/ PerServerQueue:URL;name</p> <p><b>Metric:</b> producerCount</p>	15	5	30	TRUE

<p><b>EmsServerAsyncDBSizeHigh</b> The size of the Async database, in bytes, for the EMS Server reached its maximum. <b>Index Type(s):</b> PerServer:URL <b>Metric:</b> asyncDBSize</p>	50	100	30	FALSE
<p><b>EmsServerInboundDeltaHigh</b> The number of new incoming messages for the EMS Server has reached its maximum <b>Index Type(s):</b> PerServer <b>Metric:</b> DeltainboundMessageCount</p>	60	80	30	FALSE
<p><b>EmsServerSyncDBSizeHigh</b> The size of the Sync database, in bytes, for the EMS Server reached its maximum. <b>Index Type(s):</b> PerServer:URL <b>Metric:</b> syncDBSize</p>	50	100	30	FALSE
<p><b>EmsServerConnectionCountHigh</b> Alert is triggered when the number of connections to the server reaches the specified threshold. <b>Index Type(s):</b> PerServer:URL <b>Metric:</b> connectionCount</p>	60	80	30	FALSE
<p><b>EmsServerInMsgRateHigh</b> The number of inbound messages on the server exceeded the specified threshold. <b>Index Type(s):</b> PerServer:URL <b>Metric:</b> inboundMessageRate</p>	2	80	30	FALSE
<p><b>EmsServerMemUsedHigh</b> The percent memory used on the server exceeded the specified threshold. <b>Index Type(s):</b> PerServer:URL <b>Metric:</b> messageMemoryPct</p>	60	80	30	FALSE
<p><b>EmsServerNotStarted</b> The server state is empty. The server is not started. <b>Index Type(s):</b> PerServer:URL <b>Metric:</b> NotStarted</p>	NaN	NaN	30	FALSE
<p><b>EmsServerOutMsgRateHigh</b> The number of outbound messages on the server exceeded the specified threshold. <b>Index Type(s):</b> PerServer:URL <b>Metric:</b> outboundMessageRate</p>	60	80	30	FALSE
<p><b>EmsServerPendingMsgsHigh</b> The number of pending messages in the server queue exceeded the specified threshold. <b>Index Type(s):</b> PerServer:URL <b>Metric:</b> pendingMessageCount</p>	60	80	30	FALSE

<p><b>EmsServerPendingMsgSizeHigh</b> The size, in KB, of the pending messages stored on this EMS Server reached its maximum. <b>Index Type(s):</b> PerServer:URL <b>Metric:</b> pendingMessageCount</p>	60	80	30	FALSE
<p><b>EmsServerRouteState</b> One or more routes on the server are not active. <b>Index Type(s):</b> PerServer:URL <b>Metric:</b> Alert State</p>	NaN	NaN	30	FALSE
<p><b>EmsServerStaleData</b> The server stopped receiving data. <b>Index Type(s):</b> PerServer:URL <b>Metric:</b> Expired</p>	NaN	NaN	30	FALSE
<p><b>EmsTopicConsumerIdleTimeHigh</b> The idle time of the topic consumer has reached its maximum. This alert is triggered when there is no change in the number of incoming messages for a topic within a specified period of time (in seconds). <b>Index Type(s):</b> PerTopic;PerServerTopic <b>Metric:</b> ConsumerIdleTime</p>	60	80	30	FALSE
<p><b>EmsTopicInboundDeltaHigh</b> The number of new incoming messages for the EMS Topic has reached its maximum. <b>Index Type(s):</b> PerTopic;PerServerTopic <b>Metric:</b> DeltainboundTotalMessages</p>	60	80	30	FALSE
<p><b>EmsTopicMsgLatencyHigh</b> The time, in seconds, needed to process all pending messages based on the current outbound message rate exceeded its threshold. This alert does not take into account topics with outbound messages rates equal to zero. <b>Index Type(s):</b> PerServerTopic <b>Metric:</b> messageLatency</p>	60	80	30	FALSE
<p><b>EmsTopicProviderIdleTimeHigh</b> The topic idle time exceeded the specified threshold. A topic is idle when the number of inbound messages remains unchanged. <b>Index Type(s):</b> PerServerTopic:URL;name <b>Metric:</b> ProviderIdleTime</p>	60	80	30	FALSE
<p><b>EmsTopicsConsumerCountHigh</b> The number of consumers for the topic exceeded the specified threshold. <b>Index Type(s):</b> PerServerTopic:URL;name <b>Metric:</b> consumerCount</p>	60	80	30	FALSE
<p><b>EmsTopicsConsumerCountLow</b> The number of consumers for the topic is below the specified threshold. <b>Index Type(s):</b> PerServerTopic <b>Metric:</b> consumerCount</p>	60	80	30	FALSE

<p><b>EmsTopicsInMsgRateHigh</b> The number of inbound messages for the topic exceeded the specified threshold. <b>Index Type(s):</b> PerServerTopic <b>Metric:</b> inboundMessageRate</p>	60	80	30	FALSE
<p><b>EmsTopicsOutMsgRateHigh</b> The rate of outbound messages for the topic exceeded the specified threshold. <b>Index Type(s):</b> PerServerTopic <b>Metric:</b> outboundMessageRate</p>	60	80	30	TRUE
<p><b>EmsTopicsPendingMsgsHigh</b> The number of pending messages on the queue for the topic exceeded the specified threshold. <b>Index Type(s):</b> PerTopic <b>Metric:</b> pendingMessageCount</p>	50	75	30	FALSE
<p><b>EmsTopicsProducerCountHigh</b> The number of active producers for this topic exceeded the specified high threshold. <b>Index Type(s):</b> PerTopic/PerServerTopic <b>Metric:</b> producerCount</p>	60	80	30	TRUE
<p><b>EmsTopicsProducerCountLow</b> The number of producers for the topic is below the specified threshold. <b>Index Type(s):</b> PerTopic/PerServerTopic <b>Metric:</b> producerCount</p>	60	80	30	TRUE
<p><b>EmsTopicsSubscriberCountHigh</b> The number of subscribers for the topic exceeded the specified threshold. <b>Index Type(s):</b> PerServerTopic <b>Metric:</b> subscriberCount</p>	50	75	30	FALSE
<p><b>JvmCpuPercentHigh</b> The percent JVM CPU usage exceeded the specified threshold. <b>Index Type(s):</b> PerJVM <b>Metric:</b> CpuPercent</p>	30	40	30	FALSE
<p><b>JvmGcDutyCycleHigh</b> The JVM Garbage Collection contains an item that exceeded the specified duty cycle threshold (the percent of time spent in Garbage Collection). <b>Index Type(s):</b> PerGC <b>Metric:</b> TimeUsedPercent</p>	50	75	30	FALSE
<p><b>JvmMemoryUsedHigh</b> The percent JVM memory used exceeded the specified threshold. <b>Index Type(s):</b> PerJVM <b>Metric:</b> MemoryUsedPercent</p>	50	75	30	FALSE

<b>JvmNotConnected</b> The JVM is not connected. <b>Index Type(s):</b> PerJVM <b>Metric:</b> Connected	NaN	NaN	30	FALSE
<b>JvmStaleData</b> The JVM stopped receiving data. <b>Index Type(s):</b> PerJVM <b>Metric:</b> Expired	NaN	NaN	30	FALSE

---

## TIBCO FTL

The following alerts are available for TIBCO FTL. Default settings for warning and alarm thresholds, duration and whether the alert is enabled (true/false) are shown.

Alert	Warning Level	Alarm Level	Duration	Enabled
<b>TftlClientCpuTime</b> Executes a single warning and a single alarm if the CPU response time to the client exceeds the specified threshold. Index Type: PerClient Metric: Delta_CPU_TIME	160000	200000	30	FALSE
<b>TftlClientCpuUsage</b> Executes a single warning and a single alarm if the CPU usage by the client exceeds the specified threshold. Index Type: PerClient Metric:	160000	200000	30	FALSE
<b>TftlClientExpired</b> Executes a single alert if the response time to the client exceeds the specified threshold. Index Type: PerClient Metric:	NaN	NaN	30	FALSE
<b>TftlClientMemory</b> Executes a single warning and a single alarm if the memory usage by the client exceeds the specified threshold. Index Type: PerClient Metric: PROCESS_RSS_KB	160000	200000	30	FALSE
<b>TftlClientMsgsRcvdRate</b> Executes a single warning and a single alarm if the number of messages received by the client per second exceeds the specified threshold. Index Type: PerClient Metric: RateMESSAGES_RECEIVED	160000	200000	30	FALSE

<b>TftlClientMsgsSentRate</b> Executes a single warning and a single alarm if the number of messages sent by the client per second exceeds the specified threshold. Index Type: PerClient Metric: RateMESSAGES_SENT	160000	200000	30	FALSE
<b>TftlClientNotRunning</b> Executes a single if the client status is not "RUNNING". Index Type: PerClient Metric: Delayed Writes	NaN	NaN	30	FALSE
<b>TftlClientVirtualMemory</b> Executes a single warning and a single alarm if the virtual memory usage by the client exceeds the specified threshold. Index Type: PerClient Metric:	160000	200000	30	FALSE
<b>TftlServerClientCount</b> Executes a single warning and a single alarm if the number of clients on the FTL server exceeds the specified threshold. Index Type: PerServer Metric:	160	200	30	FALSE
<b>TftlServerCpuTime</b> Executes a single warning and a single alarm if the FTL server CPU response time exceeds the specified threshold. Index Type: PerServer Metric:	160	200	30	FALSE
<b>TftlServerCpuUsage</b> Executes a single warning and a single alarm if the FTL server CPU usage exceeds the specified threshold. Index Type: PerServer Metric:	60	80	30	FALSE
<b>TftlServerExpired</b> Executes a single warning and a single alarm if the FTL server response time exceeds the specified threshold. Index Type: PerServer Metric:	NaN	NaN	30	FALSE
<b>TftlServerInboxSendFaults</b> Executes a single warning and a single alarm if the number of times the FTL server fails to queue messages to the appropriate inbox exceeds the specified threshold. Index Type: PerServer Metric: SEND_TO_INBOX_FAILURES	160	200	30	FALSE

<p><b>TftlServerMemory</b> Executes a single warning and a single alarm if the FTL server memory usage exceeds the specified threshold. Index Type: PerServer Metric: PROCESS_RSS_KB</p>	160	200	30	FALSE
<p><b>TftlServerOnBackup</b> Executes a single alert if the primary FTL server is down and now running on the backup FTL server. Index Type: PerServer Metric:</p>	NaN	NaN	30	FALSE
<p><b>TftlServerSatelliteCount</b> Executes a single alert if the number of satellite servers is lower than expected. Note: Set threshold to one less than number of deployed satellites. Index Type: Response Time Metric: Table_locks_waited</p>	NaN	5	30	FALSE
<p><b>TftlServerVirtualMemory</b> Executes a single warning and a single alarm if the FTL server virtual memory usage exceeds the specified threshold. Index Type: Response Time Metric: Table_locks_waited</p>	160	200	30	FALSE

## UX

The following are the Monitor alerts you can enable to be aware of any web application that is unresponsive, performing slowly, generating errors or returning invalid information. By default, Monitor alerts are disabled.

Monitor alerts execute when the UX Robot performs its routine runs on URLs. The **uxmon.properties** file defines which URLs the UX Robot checks and reports on. There are two types of Monitor alerts, UX-ROBOT alerts and UX-URL alerts.

- UX-ROBOT alerts apply to multiple URLs.
- UX-URL alerts apply to a single URL.

<p><b>UXRobotError</b></p>	<p>During a UX Robot run, this UX-ROBOT alert executes a single warning alert and a single alarm alert if the number of URL errors exceed the specified threshold. The warning default threshold is <b>1</b> and the alarm default threshold is <b>10</b>. For example, the URL error message "no such URL" indicates an issue at the Web Server that serves the URL. Using the default settings, a warning alert executes if the UX Robot encounters 1 or more URL errors and an alarm alert executes if the UX Robot encounters 10 or more URL errors.</p>
<p><b>UXRobotResponseSlow</b></p>	<p>During a UX Robot run, this UX-ROBOT alert executes a single warning alert and a single alarm alert if the total response time for all specified URLs exceeds the specified threshold. The warning default threshold is <b>1000</b> milliseconds and the alarm default threshold is <b>2000</b> milliseconds.</p>

<b>UXRobotSearchSentinel</b>	<p>During a UX Robot run, this UX-ROBOT alert executes a single warning alert and a single alarm alert if the number of false URL responses (responses without the specified <b>searchString</b> in the URL line) exceeds the specified threshold. The warning default threshold is <b>1</b> and the alarm default threshold is <b>10</b>.</p> <p>For example, using the default settings, a warning alert executes if the UX Robot encounters 1 or more false responses from URLs and an alarm alert executes if the UX Robot encounters 10 or more false responses from URLs.</p>
<b>UXRobotTimeout</b>	<p>During a UX Robot run, this UX-ROBOT alert executes a single warning alert and a single alarm alert if the number of URL timeouts exceeds the specified <b>maxTimeoutMS</b> threshold. The warning default threshold is <b>1</b> and the alarm default threshold is <b>15</b>.</p> <p>For example, the URL error message "no such URL" indicates an issue at the Web Server that serves the URL. Using the default settings, a warning alert executes if the UX Robot encounters 1 or more URL errors and an alarm alert executes if the UX Robot encounters 15 or more URL errors.</p>
<b>UXURLError</b>	<p>During a UX Robot run, this UX-URL alert executes a single alert if the UX Robot receives an error message from a URL. The default setting is <b>TRUE</b>.</p> <p>For example, the URL error message "no such URL" indicates an issue at the Web Server that serves the URL.</p>
<b>UXURLResponseSlow</b>	<p>During a UX Robot run, this UX-URL alert executes a single warning alert and a single alarm alert if the response time for a URL exceeds the specified threshold. The warning default threshold is <b>1000</b> milliseconds and the alarm default threshold is <b>2000</b> milliseconds.</p>
<b>UXURLSearchSentinel</b>	<p>During a UX Robot run, this UX-URL alert executes an alert if the UX Robot receives a false URL response (a response without the specified <b>searchString</b> in the URL line). The default setting is <b>FALSE</b>.</p>
<b>UXURLTimeout</b>	<p>During a UX Robot run, this UX-URL alert executes an alert if the URL response time exceeds the specified <b>maxTimeoutMS</b> threshold. UX Robot receives a false URL response (a response without the specified <b>searchString</b>). The default setting is <b>TRUE</b>.</p>

## VMware vCenter

The following alerts are available for VMware vCenter. Default settings for warning and alarm thresholds, duration and whether the alert is enabled (true/false) are shown.

Alert Name	WARN. LEVEL	ALARM LEVEL	DURATION	ENABLED
<b>VmwHostCpuUtilizationHigh</b> The Host's CPU utilization is above the defined threshold. <b>Index Type(s):</b> PerVmHost	50	75	2	TRUE
<b>VmHostDiskBytesReadHigh</b> The disk read rate (kBytes/second) is above the defined thresholds. <b>Index Type(s):</b> PerVmHost	1024	2048	2	TRUE

<b>VmHostDiskBytesWrittenHigh</b> The disk write rate (kBytes/second) is above the defined thresholds. <b>Index Type(s):</b> PerVmHost	1024	2048	2	TRUE
<b>VmwHostInBytesHigh</b> The inbound byte rate (KB/second) is above the defined thresholds. <b>Index Type(s):</b> PerVmHost	1024	2048	2	TRUE
<b>VmwHostInPktDropLossHigh</b> The percentage of inbound packets dropped is above the defined threshold. <b>Index Type(s):</b> PerVmHost	1	3	2	TRUE
<b>VmwHostInPktErrorLossHigh</b> The percentage of inbound packets discarded for any error is above the defined threshold. <b>Index Type(s):</b> PerVmHost	1	3	2	TRUE
<b>VmwHostMemoryUsageHigh</b> The percentage memory utilization (used/configured) is above the defined threshold. <b>Index Type(s):</b> PerVmHost	70	80	2	TRUE
<b>VmwHostOutBytesHigh</b> The outbound byte rate (KB/second) is above the defined threshold. <b>Index Type(s):</b> PerVmHost	1024	2048	2	TRUE
<b>VmwHostOutPktDropLossHigh</b> The percentage of outbound packets dropped is above the defined thresholds. <b>Index Type(s):</b> PerVmHost	1	3	2	TRUE
<b>VmwHostOutPktErrorLossHigh</b> The percentage of inbound packets discarded for any error is above the defined threshold. <b>Index Type(s):</b> PerVmHost	1	3	2	TRUE
<b>VmwHostStatusBad</b> The overall status is not "green." <b>Index Type(s):</b> PerVmHost	NaN	NaN	2	TRUE
<b>VmwHostSwapUsedHigh</b> The amount of swap space used by a host is above the defined thresholds. <b>Index Type(s):</b> PerVmHost	10240	40960	2	TRUE
<b>VmwVmCpuUtilizationHigh</b> The virtual machine CPU utilization is above the defined thresholds. <b>Index Type(s):</b> PerVm	50	75	2	TRUE
<b>VmwVmDiskBytesReadHigh</b> The disk read rate (KB/second) is above the defined thresholds. <b>Index Type(s):</b> PerVm	1024	2048	2	TRUE

<b>VmwVmDiskBytesWrittenHigh</b> The disk write rate (KB/second) is above the defined thresholds. <b>Index Type(s):</b> PerVm	1024	2048	2	TRUE
<b>VmwVmDiskUsageHigh</b> The amount of disk space used by the virtual machine is above the defined threshold. <b>Index Type(s):</b> PerVm	85	95	30	TRUE
<b>VmwVmInBytesHigh</b> The inbound byte rate (KB/second) is above the defined threshold. <b>Index Type(s):</b> PerVm	1024	2048	2	TRUE
<b>VmwVmInPktDropLossHigh</b> The percentage of inbound packet loss due to dropped packets is above the defined threshold. <b>Index Type(s):</b> PerVm	1	3	30	TRUE
<b>VmwVmMemoryUsageHigh</b> The percentage of memory utilization (active/ configured) is above the defined thresholds. <b>Index Type(s):</b> PerVm	70	80	2	TRUE
<b>VmwVmOutBytesHigh</b> The outbound byte rate is above the defined threshold. <b>Index Type(s):</b> PerVm	1024	2048	2	TRUE
<b>VmwVmOutPktDropLossHigh</b> The percentage of outbound packet loss due to dropped packets on the virtual machine is above the defined threshold. <b>Index Type(s):</b> PerVm	1	3	2	TRUE
<b>VmwVmStatusBad</b> The overall status for this virtual machine is not "green." <b>Index Type(s):</b> PerVm	NaN	NaN	2	TRUE
<b>VmwVmSwapUsedHigh</b> The amount of host memory swapped out for the virtual machine by the host's virtual machine kernel is above the defined threshold. This metric is not related to any swapping that may occur in the guest operating system. <b>Index Type(s):</b> PerVm	3072	4096	2	TRUE

## APPENDIX B Limitations

This section includes:

- [“iPad Safari Limitations”](#)
- [“TIBCO ActiveMatrix BusinessWorks”](#)

---

### iPad Safari Limitations

- In the iPad settings for Safari, **JavaScript** must be **ON** and **Block Pop-ups** must be **OFF**. As of this writing, the Thin Client has been tested only on iOS 4.3.5 in Safari.
- The iPad does not support Adobe Flash, so the Fx graph objects (obj\_fxtrend, obj\_fxpie, obj\_fxbar) are unavailable. The Thin Client automatically replaces the Fx graph objects with the equivalent non-Fx object (obj\_trendgraph02, obj\_pie, obj\_bargraph). Note that the replacement objects behave the same as the Fx objects in most cases but not in all. In particular, obj\_trendgraph02 does not support the sliding cursor object nor the **legendPosition** property. Custom Fx objects are not supported on the iPad.
- The Thin Client implements scrollbars for table objects and graph objects. However, unlike the scrollbars used on desktop browsers, the scrollbars used on the iPad do not have arrow buttons at each end. This can make it difficult to scroll precisely (for example, row by row) on objects with a large scrolling range.
- At full size, users may find it difficult to touch the intended display object without accidentally touching nearby objects and performing an unwanted drill-down, sort, scroll, and so forth. This is particularly true of table objects that support drill-down and also scrolling, and also in panel layouts that contain the tree navigation control. In those cases, the user may want to zoom the iPad screen before interacting with the Thin Client.
- If the iPad sleeps or auto-locks while a Thin Client display is open in Safari, or if the Safari application is minimized by clicking on the iPad's home button, the display is not updated until the iPad is awakened and Safari is reopened. In some cases it may be necessary to refresh the page from Safari's navigation bar.

Because the iPad uses a touch interface there are differences in the Thin Client appearance and behavior in iOS Safari as compared to the conventional desktop browsers that use a cursor (mouse) interface, such as Firefox and Internet Explorer. These are described below.

- **Popup browser windows:** An RTView object's drill-down target can be configured to open a display in a new window. In a desktop browser, when the RTView object is clicked the drill-down display is opened in a popup browser window. But in iOS Safari 4.3.5, only one page is visible at a time, so when the RTView object is touched a new page containing the drill-down display opens and fills the screen. The Safari navigation bar can be used to toggle between the currently open pages or close them.
- **Mouseover text:** When mouseover text and drill-down are both enabled on an RTView object (for example, a bar graph), in iOS Safari the first touch on an element in the object (for example, a bar) displays the mouseover text for that element and the second touch on the same element performs the drill-down.
- **Resize Mode and Layout:** By default, the Display Server runs with **resizeMode** set to **crop**. In **crop** mode, if a display is larger than the panel that contains it only a portion of the display is visible. In a desktop browser, scrollbars become available to allow the user to scroll to view the entire display. In iOS Safari, scrollbars do not appear but the display can be scrolled by dragging two fingers inside the display. (Dragging one finger scrolls the entire page, not the display).

If the Display Server is run with **resizeMode** set to **scale** or **layout**, the display is resized to fit into the panel that contains it. If a desktop browser is resized after a display is opened, the display is resized accordingly. On the iPad, the Safari browser can only be resized by reorienting the iPad itself, between portrait mode and landscape mode.

The panel layout feature is supported in the Thin Client. However, unlike a desktop browser which resizes to match the layout size, the size of Safari is fixed. So if the Display Server is run with **resizeMode** set to **crop** or **scale** mode, there may be unused space at the edges of the display(s) or, in **crop** mode, the panels and displays may be cropped.

This means that **layout** mode should be used for best results on the iPad. For layout mode to be most effective, displays should use the **anchor** and **dock** object properties. Please see RTView documentation for more information.

- **Scrolling:** The Thin Client implements scrollbars for table objects and graph objects. The scrollbars are activated by dragging with one finger.

If an RTView display is viewed in **crop** mode and is too large to be displayed entirely in Safari, scrollbars do not appear (as they would in a desktop browser) but the display can be scrolled by dragging with two fingers inside the display.

Scrollbars do not ever appear in a text area control. If the text area contains more text than is visible, use the two finger drag in the text area to scroll the text.

Regardless of the size of a listbox control, it can only display a single item (typically, the selected item). When the listbox is touched, the list of items appear in a popup list. In other words, on iOS Safari the listbox control and the combobox control behave identically.

- Context menu: The Thin Client context menu is opened by a right mouse button click in a desktop browser. It is opened in iOS Safari by touching any location on a display and holding that touch for 2 seconds. The menu appears in the top left corner of the display, regardless of where the display is touched. The items **Export Table to Excel**, **Drill Down**, and **Execute Command** are not included on the context menu in Safari. All other items are available. The **Export Table to HTML** item is enabled if a table object is touched (unless the table object's drillDownTarget is configured to open another display). After an **Export to PDF/HTML** is performed, the exported content opens on another page in Safari. From there, the content can either be opened by another application (for example, the iBooks application opens PDF) and emailed, or it can be copied and pasted into an email.

---

## TIBCO ActiveMatrix BusinessWorks

### Servers

#### AIX

- Status will be **LIMITED**.
- CPU Usage, Free Memory and Virtual Memory Usage will not be available.

### Business Works 5.7.1 Engine Status

The BW Engine microagent has a method **GetExecInfo** that includes a field called **Status**, which may have the following values:

- ACTIVE
- SUSPENDED
- STANDBY
- STOPPING
- STOPPED

In Business Works 5.7.1 (but not earlier or later versions) this method fails to return any data and, in some cases when the starts, it may not know an engine's exact status. For example, if an engine is deployed but not active it could be SUSPENDED or STOPPED, or else it could be ACTIVE or STOPPING. In these cases the sets the status to UNKNOWN. An UNKNOWN status will be resolved once the engine is stopped and restarted; henceforth the status will display as STOPPED or ACTIVE.

## **BWSE Components**

- JVM memory metrics are available for BWSE components running in AMX 3.x environments only.
- The BW Version column in the All Engines Table display is blank for BWSE components.
- The Deployment column in the All Engines Table display is UNKNOWN for BWSE components. This is because the AMX environment controls in which node or nodes a BWSE component is running, therefore the concept of "deployment" in traditional BusinessWorks does not apply.
- BWSE components only appear in the All Engines Table display when they are running in a node.

## APPENDIX C Third Party Notice Requirements

This section includes:

- "RTView Enterprise"
- "RTView Core"

---

### RTView Enterprise

\*\* The UX Monitor Solution Package uses HTMLUnit 2.14 and is licensed under the Apache License Version 2.0

\*\* The UX Monitor Solution Package uses Log4J 1.2 and is licensed under the Apache License Version 2.0

\*\* The RTView Host Agent Solution Package uses Hyperic SIGAR and is licensed under the Apache License Version 2.0

\*\* The Solution Package for Apache Kafka uses Jackson 2.8.4 and is licensed under the Apache License Version 2.0

\*\* The Solution Package for Apache Kafka uses Netty 3.10.5.Final and is licensed under the Apache License Version 2.0

#### Apache License

Version 2.0, January 2004

<http://www.apache.org/licenses/>

#### TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

##### 1. Definitions.

"License" shall mean the terms and conditions for use, reproduction, and distribution as defined by Sections 1 through 9 of this document.

"Licensor" shall mean the copyright owner or entity authorized by the copyright owner that is granting the License.

"Legal Entity" shall mean the union of the acting entity and all other entities that control, are controlled by, or are under common control with that entity. For the purposes of this definition, "control" means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity.

"You" (or "Your") shall mean an individual or Legal Entity exercising permissions granted by this License.

"Source" form shall mean the preferred form for making modifications, including but not limited to software source code, documentation source, and configuration files.

"Object" form shall mean any form resulting from mechanical transformation or translation of a Source form, including but not limited to compiled object code, generated documentation, and conversions to other media types.

"Work" shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in or attached to the work (an example is provided in the Appendix below)

"Derivative Works" shall mean any work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial revisions, annotations, elaborations, or other modifications represent, as a whole, an original work of authorship. For the purposes of this License, Derivative Works shall not include works that remain separable from, or merely link (or bind by name) to the interfaces of, the Work and Derivative Works thereof.

"Contribution" shall mean any work of authorship, including the original version of the Work and any modifications or additions to that Work or Derivative Works thereof, that is intentionally submitted to Licensor for inclusion in the Work by the copyright owner or by an individual or Legal Entity authorized to submit on behalf of the copyright owner. For the purposes of this definition, "submitted" means any form of electronic, verbal, or written communication sent to the Licensor or its representatives, including but not limited to communication on electronic mailing lists, source code control systems, and issue tracking systems that are managed by, or on behalf of, the Licensor for the purpose of discussing and improving the Work, but excluding communication that is conspicuously marked or otherwise designated in writing by the copyright owner as "Not a Contribution."

"Contributor" shall mean Licensor and any individual or Legal Entity on behalf of whom a Contribution has been received by Licensor and subsequently incorporated within the Work.

2. Grant of Copyright License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable copyright license to reproduce, prepare Derivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Source or Object form.

3. Grant of Patent License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable (except as stated in this section) patent license to make, have made, use, offer to sell, sell, import, and otherwise transfer the Work, where such license applies only to those patent claims licensable by such Contributor that are necessarily infringed by their Contribution(s) alone or by combination of their Contribution(s) with the Work to which such Contribution(s) was submitted. If You institute patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent infringement, then any patent licenses granted to You under this License for that Work shall terminate as of the date such litigation is filed.

4. Redistribution. You may reproduce and distribute copies of the Work or Derivative Works thereof in any medium, with or without modifications, and in Source or Object form, provided that You meet the following conditions:

(a) You must give any other recipients of the Work or Derivative Works a copy of this License; and

(b) You must cause any modified files to carry prominent notices stating that You changed the files; and

(c) You must retain, in the Source form of any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form of the Work, excluding those notices that do not pertain to any part of the Derivative Works; and

(d) If the Work includes a "NOTICE" text file as part of its distribution, then any Derivative Works that You distribute must include a readable copy of the attribution notices contained within such NOTICE file, excluding those notices that do not pertain to any part of the Derivative Works, in at least one of the following places: within a NOTICE text file distributed as part of the Derivative Works; within the Source form or documentation, if provided along with the Derivative Works; or, within a display generated by the Derivative Works, if and wherever such third-party notices normally appear. The contents of the NOTICE file are for informational purposes only and do not modify the License. You may add Your own attribution notices within Derivative Works that You distribute, alongside or as an addendum to the NOTICE text from the Work, provided that such additional attribution notices cannot be construed as modifying the License.

You may add Your own copyright statement to Your modifications and may provide additional or different license terms and conditions for use, reproduction, or distribution of Your modifications for any such Derivative Works as a whole, provided Your use, reproduction, and distribution of the Work otherwise complies with the conditions stated in this License.

5. Submission of Contributions. Unless You explicitly state otherwise, any Contribution intentionally submitted for inclusion in the Work by You to the Licensor shall be under the terms and conditions of this License, without any additional terms or conditions. Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions.

6. Trademarks. This License does not grant permission to use the trade names, trademarks, service marks, or product names of the Licensor, except as required for reasonable and customary use in describing the origin of the Work and reproducing the content of the NOTICE file.

7. Disclaimer of Warranty. Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License.

8. Limitation of Liability. In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, including any direct, indirect, special, incidental, or consequential damages of any character arising as a result of this License or out of the use or inability to use the Work (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses), even if such Contributor has been advised of the possibility of such damages.

9. Accepting Warranty or Additional Liability. While redistributing the Work or Derivative Works thereof, You may choose to offer, and charge a fee for, acceptance of support, warranty, indemnity, or other liability obligations and/or rights consistent with this License. However, in accepting such obligations, You may act only on Your own behalf and on Your sole responsibility, not on behalf of any other Contributor, and only if You agree to indemnify, defend, and hold each Contributor harmless for any liability incurred by, or claims asserted against, such Contributor by reason of your accepting any such warranty or additional liability.

## END OF TERMS AND CONDITIONS

**APPENDIX: How to apply the Apache License to your work.**

To apply the Apache License to your work, attach the following boilerplate notice, with the fields enclosed by brackets "[ ]" replaced with your own identifying information. (Don't include the brackets!) The text should be enclosed in the appropriate comment syntax for the file format. We also recommend that a file or class name and description of purpose be included on the same "printed page" as the copyright notice for easier identification within third-party archives.

Copyright [yyyy] [name of copyright owner]

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at:

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

=====  
 \*\* The UX Monitor Solution Package uses Netbeans App Framework 1.0.3 and is licensed under GNU Lesser General Public License Version 3

=====  
 \*\* BrowserLauncher2 1.3

This product uses BrowserLauncher 1.3 and is distributed pursuant to the terms of the Lesser General Public License. The source code for BrowserLauncher2 1.3 can be obtained from: <http://browserlaunch2.sourceforge.net/>

**GNU LESSER GENERAL PUBLIC LICENSE**

Version 3, 29 June 2007

Copyright © 2007 Free Software Foundation, Inc. <<http://fsf.org/>>

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

This version of the GNU Lesser General Public License incorporates the terms and conditions of version 3 of the GNU General Public License, supplemented by the additional permissions listed below.

**0. Additional Definitions.**

As used herein, "this License" refers to version 3 of the GNU Lesser General Public License, and the "GNU GPL" refers to version 3 of the GNU General Public License.

"The Library" refers to a covered work governed by this License, other than an Application or a Combined Work as defined below.

An "Application" is any work that makes use of an interface provided by the Library, but which is not otherwise based on the Library. Defining a subclass of a class defined by the Library is deemed a mode of using an interface provided by the Library.

A “Combined Work” is a work produced by combining or linking an Application with the Library. The particular version of the Library with which the Combined Work was made is also called the “Linked Version”.

The “Minimal Corresponding Source” for a Combined Work means the Corresponding Source for the Combined Work, excluding any source code for portions of the Combined Work that, considered in isolation, are based on the Application, and not on the Linked Version.

The “Corresponding Application Code” for a Combined Work means the object code and/or source code for the Application, including any data and utility programs needed for reproducing the Combined Work from the Application, but excluding the System Libraries of the Combined Work.

### **1. Exception to Section 3 of the GNU GPL.**

You may convey a covered work under sections 3 and 4 of this License without being bound by section 3 of the GNU GPL.

### **2. Conveying Modified Versions.**

If you modify a copy of the Library, and, in your modifications, a facility refers to a function or data to be supplied by an Application that uses the facility (other than as an argument passed when the facility is invoked), then you may convey a copy of the modified version:

- a) under this License, provided that you make a good faith effort to ensure that, in the event an Application does not supply the function or data, the facility still operates, and performs whatever part of its purpose remains meaningful, or
- b) under the GNU GPL, with none of the additional permissions of this License applicable to that copy.

### **3. Object Code Incorporating Material from Library Header Files.**

The object code form of an Application may incorporate material from a header file that is part of the Library. You may convey such object code under terms of your choice, provided that, if the incorporated material is not limited to numerical parameters, data structure layouts and accessors, or small macros, inline functions and templates (ten or fewer lines in length), you do both of the following:

- a) Give prominent notice with each copy of the object code that the Library is used in it and that the Library and its use are covered by this License.
- b) Accompany the object code with a copy of the GNU GPL and this license document.

### **4. Combined Works.**

You may convey a Combined Work under terms of your choice that, taken together, effectively do not restrict modification of the portions of the Library contained in the Combined Work and reverse engineering for debugging such modifications, if you also do each of the following:

- a) Give prominent notice with each copy of the Combined Work that the Library is used in it and that the Library and its use are covered by this License.
- b) Accompany the Combined Work with a copy of the GNU GPL and this license document.
- c) For a Combined Work that displays copyright notices during execution, include the copyright notice for the Library among these notices, as well as a reference directing the user to the copies of the GNU GPL and this license document.
- d) Do one of the following:

- 0) Convey the Minimal Corresponding Source under the terms of this License, and the Corresponding Application Code in a form suitable for, and under terms that permit, the user to recombine or relink the Application with a modified version of the Linked Version to produce a modified Combined Work, in the manner specified by section 6 of the GNU GPL for conveying Corresponding Source.
- 1) Use a suitable shared library mechanism for linking with the Library. A suitable mechanism is one that (a) uses at run time a copy of the Library already present on the user's computer system, and (b) will operate properly with a modified version of the Library that is interface-compatible with the Linked Version.
- e) Provide Installation Information, but only if you would otherwise be required to provide such information under section 6 of the GNU GPL, and only to the extent that such information is necessary to install and execute a modified version of the Combined Work produced by recombining or relinking the Application with a modified version of the Linked Version. (If you use option 4d0, the Installation Information must accompany the Minimal Corresponding Source and Corresponding Application Code. If you use option 4d1, you must provide the Installation Information in the manner specified by section 6 of the GNU GPL for conveying Corresponding Source.)

### **5. Combined Libraries.**

You may place library facilities that are a work based on the Library side by side in a single library together with other library facilities that are not Applications and are not covered by this License, and convey such a combined library under terms of your choice, if you do both of the following:

- a) Accompany the combined library with a copy of the same work based on the Library, uncombined with any other library facilities, conveyed under the terms of this License.
- b) Give prominent notice with the combined library that part of it is a work based on the Library, and explaining where to find the accompanying uncombined form of the same work.

### **6. Revised Versions of the GNU Lesser General Public License.**

The Free Software Foundation may publish revised and/or new versions of the GNU Lesser General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Library as you received it specifies that a certain numbered version of the GNU Lesser General Public License "or any later version" applies to it, you have the option of following the terms and conditions either of that published version or of any later version published by the Free Software Foundation. If the Library as you received it does not specify a version number of the GNU Lesser General Public License, you may choose any version of the GNU Lesser General Public License ever published by the Free Software Foundation.

If the Library as you received it specifies that a proxy can decide whether future versions of the GNU Lesser General Public License shall apply, that proxy's public statement of acceptance of any version is permanent authorization for you to choose that version for the Library.

---

## RTView Core

\*\* Apache Tomcat is delivered for convenience only as a separate application and is licensed under the Apache License Version 2.0

\*\* Apache HttpClient is embedded in the RTView Core libraries and is licensed under the Apache License Version 2.0

\*\* JEval 0.9.4 is licensed under the Apache License Version 2.0

\*\* Jetty 9.2.19 is licensed under the Apache License Version 2.0

### Apache License

Version 2.0, January 2004

<http://www.apache.org/licenses/>

### TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

#### 1. Definitions.

"License" shall mean the terms and conditions for use, reproduction, and distribution as defined by Sections 1 through 9 of this document.

"Licensor" shall mean the copyright owner or entity authorized by the copyright owner that is granting the License.

"Legal Entity" shall mean the union of the acting entity and all other entities that control, are controlled by, or are under common control with that entity. For the purposes of this definition, "control" means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity.

"You" (or "Your") shall mean an individual or Legal Entity exercising permissions granted by this License.

"Source" form shall mean the preferred form for making modifications, including but not limited to software source code, documentation source, and configuration files.

"Object" form shall mean anyform resulting from mechanical transformation or translation of a Source form, including but not limited to compiled object code, generated documentation, and conversions to other media types.

"Work" shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in or attached to the work (an example is provided in the Appendix below)

"Derivative Works" shall mean any work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial revisions, annotations, elaborations, or other modifications represent, as a whole, an original work of authorship. For the purposes of this License, Derivative Works shall not include works that remain separable from, or merely link (or bind by name) to the interfaces of, the Work and Derivative Works thereof.

"Contribution" shall mean any work of authorship, including the original version of the Work and any modifications or additions to that Work or Derivative Works thereof, that is intentionally submitted to Licensor for inclusion in the Work by the copyright owner or by an individual or Legal Entity authorized to submit on behalf of the copyright owner. For the purposes of this definition, "submitted" means any form of electronic, verbal, or written communication sent to the Licensor or its representatives, including but not limited to communication on electronic mailing lists, source code control systems, and issue tracking systems that are managed by, or on behalf of, the Licensor for the purpose of discussing and improving the Work, but excluding communication that is conspicuously marked or otherwise designated in writing by the copyright owner as "Not a Contribution."

"Contributor" shall mean Licensor and any individual or Legal Entity on behalf of whom a Contribution has been received by Licensor and subsequently incorporated within the Work.

2. Grant of Copyright License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable copyright license to reproduce, prepare Derivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Source or Object form.

3. Grant of Patent License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable (except as stated in this section) patent license to make, have made, use, offer to sell, sell, import, and otherwise transfer the Work, where such license applies only to those patent claims licensable by such Contributor that are necessarily infringed by their Contribution(s) alone or by combination of their Contribution(s) with the Work to which such Contribution(s) was submitted. If You institute patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent infringement, then any patent licenses granted to You under this License for that Work shall terminate as of the date such litigation is filed.

4. Redistribution. You may reproduce and distribute copies of the Work or Derivative Works thereof in any medium, with or without modifications, and in Source or Object form, provided that You meet the following conditions:

- (a) You must give any other recipients of the Work or Derivative Works a copy of this License; and
- (b) You must cause any modified files to carry prominent notices stating that You changed the files; and
- (c) You must retain, in the Source form of any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form of the Work, excluding those notices that do not pertain to any part of the Derivative Works; and
- (d) If the Work includes a "NOTICE" text file as part of its distribution, then any Derivative Works that You distribute must include a readable copy of the attribution notices contained within such NOTICE file, excluding those notices that do not pertain to any part of the Derivative Works, in at least one of the following places: within a NOTICE text file distributed as part of the Derivative Works; within the Source form or documentation, if provided along with the Derivative Works; or, within a display generated by the Derivative Works, if and wherever such third-party notices normally appear. The contents of the NOTICE file are for informational purposes only and do not modify the License. You may add Your own attribution notices within Derivative Works that You distribute, alongside or as an addendum to the NOTICE text from the Work, provided that such additional attribution notices cannot be construed as modifying the License.

You may add Your own copyright statement to Your modifications and may provide additional or different license terms and conditions for use, reproduction, or distribution of Your modifications for any such Derivative Works as a whole, provided Your use, reproduction, and distribution of the Work otherwise complies with the conditions stated in this License.

5. **Submission of Contributions.** Unless You explicitly state otherwise, any Contribution intentionally submitted for inclusion in the Work by You to the Licensor shall be under the terms and conditions of this License, without any additional terms or conditions. Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions.
6. **Trademarks.** This License does not grant permission to use the trade names, trademarks, service marks, or product names of the Licensor, except as required for reasonable and customary use in describing the origin of the Work and reproducing the content of the NOTICE file.
7. **Disclaimer of Warranty.** Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License.
8. **Limitation of Liability.** In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, including any direct, indirect, special, incidental, or consequential damages of any character arising as a result of this License or out of the use or inability to use the Work (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses), even if such Contributor has been advised of the possibility of such damages.
9. **Accepting Warranty or Additional Liability.** While redistributing the Work or Derivative Works thereof, You may choose to offer, and charge a fee for, acceptance of support, warranty, indemnity, or other liability obligations and/or rights consistent with this License. However, in accepting such obligations, You may act only on Your own behalf and on Your sole responsibility, not on behalf of any other Contributor, and only if You agree to indemnify, defend, and hold each Contributor harmless for any liability incurred by, or claims asserted against, such Contributor by reason of your accepting any such warranty or additional liability.

END OF TERMS AND CONDITIONS

#### **APPENDIX: How to apply the Apache License to your work.**

To apply the Apache License to your work, attach the following boilerplate notice, with the fields enclosed by brackets "[ ]" replaced with your own identifying information. (Don't include the brackets!) The text should be enclosed in the appropriate comment syntax for the file format. We also recommend that a file or class name and description of purpose be included on the same "printed page" as the copyright notice for easier identification within third-party archives.

Copyright [yyyy] [name of copyright owner]

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at:

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

=====

\*\* TreeMap Algorithms v1.0 is used without modifications and licensed by MPL Version 1.1. The source for TreeMap Algorithms can be obtained from <http://www.cs.umd.edu/hcil/treemap/>

\*\* iTextAsian 1.0 is licensed by MPL Version 1.1 and the source can be obtained from: <http://itextpdf.com/download.php>

## **MOZILLA PUBLIC LICENSE**

Version 1.1

### 1. Definitions.

1.0.1. "Commercial Use" means distribution or otherwise making the Covered Code available to a third party.

1.1. "Contributor" means each entity that creates or contributes to the creation of Modifications.

1.2. "Contributor Version" means the combination of the Original Code, prior Modifications used by a Contributor, and the Modifications made by that particular Contributor.

1.3. "Covered Code" means the Original Code or Modifications or the combination of the Original Code and Modifications, in each case including portions thereof.

1.4. "Electronic Distribution Mechanism" means a mechanism generally accepted in the software development community for the electronic transfer of data.

1.5. "Executable" means Covered Code in any form other than Source Code.

1.6. "Initial Developer" means the individual or entity identified as the Initial Developer in the Source Code notice required by Exhibit A.

1.7. "Larger Work" means a work which combines Covered Code or portions thereof with code not governed by the terms of this License.

1.8. "License" means this document.

1.8.1. "Licensable" means having the right to grant, to the maximum extent possible, whether at the time of the initial grant or subsequently acquired, any and all of the rights conveyed herein.

1.9. "Modifications" means any addition to or deletion from the substance or structure of either the Original Code or any previous Modifications. When Covered Code is released as a series of files, a Modification is:

A. Any addition to or deletion from the contents of a file containing Original Code or previous Modifications.

B. Any new file that contains any part of the Original Code or previous Modifications.

1.10. "Original Code" means Source Code of computer software code which is described in the Source Code notice required by Exhibit A as Original Code, and which, at the time of its release under this License is not already Covered Code governed by this License.

1.10.1. "Patent Claims" means any patent claim(s), now owned or hereafter acquired, including without limitation, method, process, and apparatus claims, in any patent Licensable by grantor.

1.11. "Source Code" means the preferred form of the Covered Code for making modifications to it, including all modules it contains, plus any associated interface definition files, scripts used to control compilation and installation of an Executable, or source code differential comparisons against either the Original Code or another well known, available Covered Code of the Contributor's choice. The Source Code can be in a compressed or archival form, provided the appropriate decompression or de-archiving software is widely available for no charge.

1.12. "You" (or "Your") means an individual or a legal entity exercising rights under, and complying with all of the terms of, this License or a future version of this License issued under Section 6.1. For legal entities, "You" includes any entity which controls, is controlled by, or is under common control with You. For purposes of this definition, "control" means (a) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (b) ownership of more than fifty percent (50%) of the outstanding shares or beneficial ownership of such entity.

## 2. Source Code License.

### 2.1. The Initial Developer Grant.

The Initial Developer hereby grants You a world-wide, royalty-free, non-exclusive license, subject to third party intellectual property claims:

(a) under intellectual property rights (other than patent or trademark) Licensable by Initial Developer to use, reproduce, modify, display, perform, sublicense and distribute the Original Code (or portions thereof) with or without Modifications, and/or as part of a Larger Work; and

(b) under Patent Claims infringed by the making, using or selling of Original Code, to make, have made, use, practice, sell, and offer for sale, and/or otherwise dispose of the Original Code (or portions thereof).

(c) the licenses granted in this Section 2.1(a) and (b) are effective on the date Initial Developer first distributes Original Code under the terms of this License.

(d) Notwithstanding Section 2.1(b) above, no patent license is granted: 1) for code that You delete from the Original Code; 2) separate from the Original Code; or 3) for infringements caused by: i) the modification of the Original Code or ii) the combination of the Original Code with other software or devices.

### 2.2. Contributor Grant.

Subject to third party intellectual property claims, each Contributor hereby grants You a world-wide, royalty-free, non-exclusive license

(a) under intellectual property rights (other than patent or trademark) Licensable by Contributor, to use, reproduce, modify, display, perform, sublicense and distribute the Modifications created by such Contributor (or portions thereof) either on an unmodified basis, with other Modifications, as Covered Code and/or as part of a Larger Work; and

(b) under Patent Claims infringed by the making, using, or selling of Modifications made by that Contributor either alone and/or in combination with its Contributor Version (or portions of such combination), to make, use, sell, offer for sale, have made, and/or otherwise dispose of: 1) Modifications made by that Contributor (or portions thereof); and 2) the combination of Modifications made by that Contributor with its Contributor Version (or portions of such combination).

(c) the licenses granted in Sections 2.2(a) and 2.2(b) are effective on the date Contributor first makes Commercial Use of the Covered Code.

(d) Notwithstanding Section 2.2(b) above, no patent license is granted: 1) for any code that Contributor has deleted from the Contributor Version; 2) separate from the Contributor Version; 3) for infringements caused by: i) third party modifications of Contributor Version or ii) the combination of Modifications made by that Contributor with other software (except as part of the Contributor Version) or other devices; or 4) under Patent Claims infringed by Covered Code in the absence of Modifications made by that Contributor.

### 3. Distribution Obligations.

#### 3.1. Application of License.

The Modifications which You create or to which You contribute are governed by the terms of this License, including without limitation Section 2.2. The Source Code version of Covered Code may be distributed only under the terms of this License or a future version of this License released under Section 6.1, and You must include a copy of this License with every copy of the Source Code You distribute. You may not offer or impose any terms on any Source Code version that alters or restricts the applicable version of this License or the recipients' rights hereunder. However, You may include an additional document offering the additional rights described in Section 3.5.

#### 3.2. Availability of Source Code.

Any Modification which You create or to which You contribute must be made available in Source Code form under the terms of this License either on the same media as an Executable version or via an accepted Electronic Distribution Mechanism to anyone to whom you made an Executable version available; and if made available via Electronic Distribution Mechanism, must remain available for at least twelve (12) months after the date it initially became available, or at least six (6) months after a subsequent version of that particular Modification has been made available to such recipients. You are responsible for ensuring that the Source Code version remains available even if the Electronic Distribution Mechanism is maintained by a third party.

#### 3.3. Description of Modifications.

You must cause all Covered Code to which You contribute to contain a file documenting the changes You made to create that Covered Code and the date of any change. You must include a prominent statement that the Modification is derived, directly or indirectly, from Original Code provided by the Initial Developer and including the name of the Initial Developer in (a) the Source Code, and (b) in any notice in an Executable version or related documentation in which You describe the origin or ownership of the Covered Code.

#### 3.4. Intellectual Property Matters

##### (a) Third Party Claims.

If Contributor has knowledge that a license under a third party's intellectual property rights is required to exercise the rights granted by such Contributor under Sections 2.1 or 2.2, Contributor must include a text file with the Source Code distribution titled "LEGAL" which describes the claim and the party making the claim in sufficient detail that a recipient will know whom to contact. If Contributor obtains such knowledge after the Modification is made available as described in Section 3.2, Contributor shall promptly modify the LEGAL file in all copies Contributor makes available thereafter and shall take other steps (such as notifying appropriate mailing lists or newsgroups) reasonably calculated to inform those who received the Covered Code that new knowledge has been obtained.

##### (b) Contributor APIs.

If Contributor's Modifications include an application programming interface and Contributor has knowledge of patent licenses which are reasonably necessary to implement that API, Contributor must also include this information in the LEGAL file.

(c) Representations.

Contributor represents that, except as disclosed pursuant to Section 3.4(a) above, Contributor believes that Contributor's Modifications are Contributor's original creation(s) and/or Contributor has sufficient rights to grant the rights conveyed by this License.

### 3.5. Required Notices.

You must duplicate the notice in Exhibit A in each file of the Source Code. If it is not possible to put such notice in a particular Source Code file due to its structure, then You must include such notice in a location (such as a relevant directory) where a user would be likely to look for such a notice. If You created one or more Modification(s) You may add your name as a Contributor to the notice described in Exhibit A. You must also duplicate this License in any documentation for the Source Code where You describe recipients' rights or ownership rights relating to Covered Code. You may choose to offer, and to charge a fee for, warranty, support, indemnity or liability obligations to one or more recipients of Covered Code. However, You may do so only on Your own behalf, and not on behalf of the Initial Developer or any Contributor. You must make it absolutely clear that any such warranty, support, indemnity or liability obligation is offered by You alone, and You hereby agree to indemnify the Initial Developer and every Contributor for any liability incurred by the Initial Developer or such Contributor as a result of warranty, support, indemnity or liability terms You offer.

### 3.6. Distribution of Executable Versions.

You may distribute Covered Code in Executable form only if the requirements of Section 3.1-3.5 have been met for that Covered Code, and if You include a notice stating that the Source Code version of the Covered Code is available under the terms of this License, including a description of how and where You have fulfilled the obligations of Section 3.2. The notice must be conspicuously included in any notice in an Executable version, related documentation or collateral in which You describe recipients' rights relating to the Covered Code. You may distribute the Executable version of Covered Code or ownership rights under a license of Your choice, which may contain terms different from this License, provided that You are in compliance with the terms of this License and that the license for the Executable version does not attempt to limit or alter the recipient's rights in the Source Code version from the rights set forth in this License. If You distribute the Executable version under a different license You must make it absolutely clear that any terms which differ from this License are offered by You alone, not by the Initial Developer or any Contributor. You hereby agree to indemnify the Initial Developer and every Contributor for any liability incurred by the Initial Developer or such Contributor as a result of any such terms You offer.

### 3.7. Larger Works.

You may create a Larger Work by combining Covered Code with other code not governed by the terms of this License and distribute the Larger Work as a single product. In such a case, You must make sure the requirements of this License are fulfilled for the Covered Code.

## 4. Inability to Comply Due to Statute or Regulation.

If it is impossible for You to comply with any of the terms of this License with respect to some or all of the Covered Code due to statute, judicial order, or regulation then You must: (a) comply with the terms of this License to the maximum extent possible; and (b) describe the limitations and the code they affect. Such description must be included in the LEGAL file described in Section 3.4 and must be included with all distributions of the Source Code. Except to the extent prohibited by statute or regulation, such description must be sufficiently detailed for a recipient of ordinary skill to be able to understand it.

## 5. Application of this License.

This License applies to code to which the Initial Developer has attached the notice in Exhibit A and to related Covered Code.

## 6. Versions of the License.

### 6.1. New Versions.

Netscape Communications Corporation ("Netscape") may publish revised and/or new versions of the License from time to time. Each version will be given a distinguishing version number.

### 6.2. Effect of New Versions.

Once Covered Code has been published under a particular version of the License, You may always continue to use it under the terms of that version. You may also choose to use such Covered Code under the terms of any subsequent version of the License published by Netscape. No one other than Netscape has the right to modify the terms applicable to Covered Code created under this License.

### 6.3. Derivative Works.

If You create or use a modified version of this License (which you may only do in order to apply it to code which is not already Covered Code governed by this License), You must (a) rename Your license so that the phrases "Mozilla", "MOZILLAPL", "MOZPL", "Netscape", "MPL", "NPL" or any confusingly similar phrase do not appear in your license (except to note that your license differs from this License) and (b) otherwise make it clear that Your version of the license contains terms which differ from the Mozilla Public License and Netscape Public License. (Filling in the name of the Initial Developer, Original Code or Contributor in the notice described in Exhibit A shall not of themselves be deemed to be modifications of this License.)

## 7. DISCLAIMER OF WARRANTY.

COVERED CODE IS PROVIDED UNDER THIS LICENSE ON AN "AS IS" BASIS, WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, WARRANTIES THAT THE COVERED CODE IS FREE OF DEFECTS, MERCHANTABILITY, FIT FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE COVERED CODE IS WITH YOU. SHOULD ANY COVERED CODE PROVE DEFECTIVE IN ANY RESPECT, YOU (NOT THE INITIAL DEVELOPER OR ANY OTHER CONTRIBUTOR) ASSUME THE COST OF ANY NECESSARY SERVICING, REPAIR OR CORRECTION. THIS DISCLAIMER OF WARRANTY CONSTITUTES AN ESSENTIAL PART OF THIS LICENSE. NO USE OF ANY COVERED CODE IS AUTHORIZED HEREUNDER EXCEPT UNDER THIS DISCLAIMER.

## 8. TERMINATION.

8.1. This License and the rights granted hereunder will terminate automatically if You fail to comply with terms herein and fail to cure such breach within 30 days of becoming aware of the breach. All sublicenses to the Covered Code which are properly granted shall survive any termination of this License. Provisions which, by their nature, must remain in effect beyond the termination of this License shall survive.

8.2. If You initiate litigation by asserting a patent infringement claim (excluding declaratory judgment actions) against Initial Developer or a Contributor (the Initial Developer or Contributor against whom You file such action is referred to as "Participant") alleging that:

(a) such Participant's Contributor Version directly or indirectly infringes any patent, then any and all rights granted by such Participant to You under Sections 2.1 and/or 2.2 of this License shall, upon 60 days notice from Participant terminate prospectively, unless if within 60 days after receipt of notice You either: (I) agree in writing to pay Participant a mutually agreeable reasonable royalty for Your past and future use of Modifications made by such Participant, or (ii) withdraw Your litigation claim with respect to the Contributor Version against such Participant. If within 60 days of notice, a reasonable royalty and payment arrangement are not mutually agreed upon in writing by the parties or the litigation claim is not withdrawn, the rights granted by Participant to You under Sections 2.1 and/or 2.2 automatically terminate at the expiration of the 60 day notice period specified above.

(b) any software, hardware, or device, other than such Participant's Contributor Version, directly or indirectly infringes any patent, then any rights granted to You by such Participant under Sections 2.1(b) and 2.2(b) are revoked effective as of the date You first made, used, sold, distributed, or had made, Modifications made by that Participant.

8.3. If You assert a patent infringement claim against Participant alleging that such Participant's Contributor Version directly or indirectly infringes any patent where such claim is resolved (such as by license or settlement) prior to the initiation of patent infringement litigation, then the reasonable value of the licenses granted by such Participant under Sections 2.1 or 2.2 shall be taken into account in determining the amount or value of any payment or license.

8.4. In the event of termination under Sections 8.1 or 8.2 above, all end user license agreements (excluding distributors and resellers) which have been validly granted by You or any distributor hereunder prior to termination shall survive termination.

## 9. LIMITATION OF LIABILITY.

UNDER NO CIRCUMSTANCES AND UNDER NO LEGAL THEORY, WHETHER TORT (INCLUDING NEGLIGENCE), CONTRACT, OR OTHERWISE, SHALL YOU, THE INITIAL DEVELOPER, ANY OTHER CONTRIBUTOR, OR ANY DISTRIBUTOR OF COVERED CODE, OR ANY SUPPLIER OF ANY OF SUCH PARTIES, BE LIABLE TO ANY PERSON FOR ANY INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES OF ANY CHARACTER INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF GOODWILL, WORK STOPPAGE, COMPUTER FAILURE OR MALFUNCTION, OR ANY AND ALL OTHER COMMERCIAL DAMAGES OR LOSSES, EVEN IF SUCH PARTY SHALL HAVE BEEN INFORMED OF THE POSSIBILITY OF SUCH DAMAGES. THIS LIMITATION OF LIABILITY SHALL NOT APPLY TO LIABILITY FOR DEATH OR PERSONAL INJURY RESULTING FROM SUCH PARTY'S NEGLIGENCE TO THE EXTENT APPLICABLE LAW PROHIBITS SUCH LIMITATION. SOME JURISDICTIONS DO NOT ALLOW THE EXCLUSION OR LIMITATION OF INCIDENTAL OR CONSEQUENTIAL DAMAGES, SO THIS EXCLUSION AND LIMITATION MAY NOT APPLY TO YOU.

## 10. U.S. GOVERNMENT END USERS.

The Covered Code is a "commercial item," as that term is defined in 48 C.F.R. 2.101 (Oct. 1995), consisting of "commercial computer software" and "commercial computer software documentation," as such terms are used in 48 C.F.R. 12.212 (Sept. 1995). Consistent with 48 C.F.R. 12.212 and 48 C.F.R. 227.7202-1 through 227.7202-4 (June 1995), all U.S. Government End Users acquire Covered Code with only those rights set forth herein.

## 11. MISCELLANEOUS.

This License represents the complete agreement concerning subject matter hereof. If any provision of this License is held to be unenforceable, such provision shall be reformed only to the extent necessary to make it enforceable. This License shall be governed by California law provisions (except to the extent applicable law, if any, provides otherwise), excluding its conflict-of-law provisions. With respect to disputes in which at least one party is a citizen of, or an entity chartered or registered to do business in the United States of America, any litigation relating to this License shall be subject to the jurisdiction of the Federal Courts of the Northern District of California, with venue lying in Santa Clara County, California, with the losing party responsible for costs, including without limitation, court costs and reasonable attorneys' fees and expenses. The application of the United Nations Convention on Contracts for the International Sale of Goods is expressly excluded. Any law or regulation which provides that the language of a contract shall be construed against the drafter shall not apply to this License.

12. RESPONSIBILITY FOR CLAIMS.

As between Initial Developer and the Contributors, each party is responsible for claims and damages arising, directly or indirectly, out of its utilization of rights under this License and You agree to work with Initial Developer and Contributors to distribute such responsibility on an equitable basis. Nothing herein is intended or shall be deemed to constitute any admission of liability.

13. MULTIPLE-LICENSED CODE.

Initial Developer may designate portions of the Covered Code as "Multiple-Licensed". "Multiple-Licensed" means that the Initial Developer permits you to utilize portions of the Covered Code under Your choice of the NPL or the alternative licenses, if any, specified by the Initial Developer in the file described in Exhibit A.

EXHIBIT A -Mozilla Public License.

The contents of this file are subject to the Mozilla Public License Version 1.1 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at <http://www.mozilla.org/MPL/>

Software distributed under the License is distributed on an "AS IS" basis, WITHOUT WARRANTY OF ANY KIND, either express or implied. See the License for the specific language governing rights and limitations under the License.

The Original Code is \_\_\_\_\_.

The Initial Developer of the Original Code is \_\_\_\_\_.

Portions created by \_\_\_\_\_ are Copyright (C) \_\_\_\_\_  
\_\_\_\_\_. All Rights Reserved.

Contributor(s): \_\_\_\_\_.

Alternatively, the contents of this file may be used under the terms of the \_\_\_\_\_ license (the "[\_\_\_\_\_] License"), in which case the provisions of [\_\_\_\_\_] License are applicable instead of those above. If you wish to allow use of your version of this file only under the terms of the [\_\_\_\_\_] License and not to allow others to use your version of this file under the MPL, indicate your decision by deleting the provisions above and replace them with the notice and other provisions required by the [\_\_\_\_\_] License. If you do not delete the provisions above, a recipient may use your version of this file under either the MPL or the [\_\_\_\_\_] License."

[NOTE: The text of this Exhibit A may differ slightly from the text of the notices in the Source Code files of the Original Code. You should use the text of this Exhibit A rather than the text found in the Original Code Source Code for Your Modifications.]

=====

**\*\*MD Datejs**

Copyright © 2006-2010 Coolite Inc.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

=====

**\*\*jQuery**

Copyright © 2009 John Resig

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

=====

**\*\* JCalendar 1.3.2**

This product uses JCalendar 1.3.2. JCalendar is distributed pursuant to the terms of the Lesser General Public License. The source code for the JCalendar may be obtained from <http://www.toedter.com/en/jcalendar/index.html>

=====

**\*\* BrowserLauncher2 1.3**

This product uses BrowserLauncher 1.3 and is distributed pursuant to the terms of the Lesser General Public License. The source code for BrowserLauncher2 1.3 can be obtained from: <http://browserlaunch2.sourceforge.net/>

**GNU LESSER GENERAL PUBLIC LICENSE**

Version 2.1, February 1999

Copyright (C) 1991, 1999 Free Software Foundation, Inc.

59 Temple Place, Suite 330, Boston, MA 02111-1307 USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

[This is the first released version of the Lesser GPL. It also counts as the successor of the GNU Library Public License, version 2, hence the version number 2.1.]

## **Preamble**

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public Licenses are intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users.

This license, the Lesser General Public License, applies to some specially designated software packages--typically libraries--of the Free Software Foundation and other authors who decide to use it. You can use it too, but we suggest you first think carefully about whether this license or the ordinary General Public License is the better strategy to use in any particular case, based on the explanations below.

When we speak of free software, we are referring to freedom of use, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish); that you receive source code or can get it if you want it; that you can change the software and use pieces of it in new free programs; and that you are informed that you can do these things.

To protect your rights, we need to make restrictions that forbid distributors to deny you these rights or to ask you to surrender these rights. These restrictions translate to certain responsibilities for you if you distribute copies of the library or if you modify it.

For example, if you distribute copies of the library, whether gratis or for a fee, you must give the recipients all the rights that we gave you. You must make sure that they, too, receive or can get the source code. If you link other code with the library, you must provide complete object files to the recipients, so that they can relink them with the library after making changes to the library and recompiling it. And you must show them these terms so they know their rights.

We protect your rights with a two-step method: (1) we copyright the library, and (2) we offer you this license, which gives you legal permission to copy, distribute and/or modify the library.

To protect each distributor, we want to make it very clear that there is no warranty for the free library. Also, if the library is modified by someone else and passed on, the recipients should know that what they have is not the original version, so that the original author's reputation will not be affected by problems that might be introduced by others.

Finally, software patents pose a constant threat to the existence of any free program. We wish to make sure that a company cannot effectively restrict the users of a free program by obtaining a restrictive license from a patent holder. Therefore, we insist that any patent license obtained for a version of the library must be consistent with the full freedom of use specified in this license.

Most GNU software, including some libraries, is covered by the ordinary GNU General Public License. This license, the GNU Lesser General Public License, applies to certain designated libraries, and is quite different from the ordinary General Public License. We use this license for certain libraries in order to permit linking those libraries into non-free programs.

When a program is linked with a library, whether statically or using a shared library, the combination of the two is legally speaking a combined work, a derivative of the original library. The ordinary General Public License therefore permits such linking only if the entire combination fits its criteria of freedom. The Lesser General Public License permits more lax criteria for linking other code with the library.

We call this license the "Lesser" General Public License because it does Less to protect the user's freedom than the ordinary General Public License. It also provides other free software developers Less of an advantage over competing non-free programs. These disadvantages are the reason we use the ordinary General Public License for many libraries. However, the Lesser license provides advantages in certain special circumstances.

For example, on rare occasions, there may be a special need to encourage the widest possible use of a certain library, so that it becomes a de-facto standard. To achieve this, non-free programs must be allowed to use the library. A more frequent case is that a free library does the same job as widely used non-free libraries. In this case, there is little to gain by limiting the free library to free software only, so we use the Lesser General Public License.

In other cases, permission to use a particular library in non-free programs enables a greater number of people to use a large body of free software. For example, permission to use the GNU C Library in non-free programs enables many more people to use the whole GNU operating system, as well as its variant, the GNU/Linux operating system.

Although the Lesser General Public License is Less protective of the users' freedom, it does ensure that the user of a program that is linked with the Library has the freedom and the wherewithal to run that program using a modified version of the Library.

The precise terms and conditions for copying, distribution and modification follow. Pay close attention to the difference between a "work based on the library" and a "work that uses the library". The former contains code derived from the library, whereas the latter must be combined with the library in order to run.

## **TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION**

0. This License Agreement applies to any software library or other program which contains a notice placed by the copyright holder or other authorized party saying it may be distributed under the terms of this Lesser General Public License (also called "this License"). Each licensee is addressed as "you".

A "library" means a collection of software functions and/or data prepared so as to be conveniently linked with application programs (which use some of those functions and data) to form executables.

The "Library", below, refers to any such software library or work which has been distributed under these terms. A "work based on the Library" means either the Library or any derivative work under copyright law: that is to say, a work containing the Library or a portion of it, either verbatim or with modifications and/or translated straightforwardly into another language. (Hereinafter, translation is included without limitation in the term "modification".)

"Source code" for a work means the preferred form of the work for making modifications to it. For a library, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the library.

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running a program using the Library is not restricted, and output from such a program is covered only if its contents constitute a work based on the Library (independent of the use of the Library in a tool for writing it). Whether that is true depends on what the Library does and what the program that uses the Library does.

1. You may copy and distribute verbatim copies of the Library's complete source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and distribute a copy of this License along with the Library.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Library or any portion of it, thus forming a work based on the Library, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

- a) The modified work must itself be a software library.
- b) You must cause the files modified to carry prominent notices stating that you changed the files and the date of any change.
- c) You must cause the whole of the work to be licensed at no charge to all third parties under the terms of this License.
- d) If a facility in the modified Library refers to a function or a table of data to be supplied by an application program that uses the facility, other than as an argument passed when the facility is invoked, then you must make a good faith effort to ensure that, in the event an application does not supply such function or table, the facility still operates, and performs whatever part of its purpose remains meaningful.

(For example, a function in a library to compute square roots has a purpose that is entirely well-defined independent of the application. Therefore, Subsection 2d requires that any application-supplied function or table used by this function must be optional: if the application does not supply it, the square root function must still compute square roots.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Library, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Library, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Library.

In addition, mere aggregation of another work not based on the Library with the Library (or with a work based on the Library) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may opt to apply the terms of the ordinary GNU General Public License instead of this License to a given copy of the Library. To do this, you must alter all the notices that refer to this License, so that they refer to the ordinary GNU General Public License, version 2, instead of to this License. (If a newer version than version 2 of the ordinary GNU General Public License has appeared, then you can specify that version instead if you wish.) Do not make any other change in these notices.

Once this change is made in a given copy, it is irreversible for that copy, so the ordinary GNU General Public License applies to all subsequent copies and derivative works made from that copy.

This option is useful when you wish to copy part of the code of the Library into a program that is not a library.

4. You may copy and distribute the Library (or a portion or derivative of it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange.

If distribution of object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place satisfies the requirement to distribute the source code, even though third parties are not compelled to copy the source along with the object code.

5. A program that contains no derivative of any portion of the Library, but is designed to work with the Library by being compiled or linked with it, is called a "work that uses the Library". Such a work, in isolation, is not a derivative work of the Library, and therefore falls outside the scope of this License.

However, linking a "work that uses the Library" with the Library creates an executable that is a derivative of the Library (because it contains portions of the Library), rather than a "work that uses the library". The executable is therefore covered by this License. Section 6 states terms for distribution of such executables.

When a "work that uses the Library" uses material from a header file that is part of the Library, the object code for the work may be a derivative work of the Library even though the source code is not. Whether this is true is especially significant if the work can be linked without the Library, or if the work is itself a library. The threshold for this to be true is not precisely defined by law.

If such an object file uses only numerical parameters, data structure layouts and accessors, and small macros and small inline functions (ten lines or less in length), then the use of the object file is unrestricted, regardless of whether it is legally a derivative work. (Executables containing this object code plus portions of the Library will still fall under Section 6.)

Otherwise, if the work is a derivative of the Library, you may distribute the object code for the work under the terms of Section 6. Any executables containing that work also fall under Section 6, whether or not they are linked directly with the Library itself.

6. As an exception to the Sections above, you may also combine or link a "work that uses the Library" with the Library to produce a work containing portions of the Library, and distribute that work under terms of your choice, provided that the terms permit modification of the work for the customer's own use and reverse engineering for debugging such modifications.

You must give prominent notice with each copy of the work that the Library is used in it and that the Library and its use are covered by this License. You must supply a copy of this License. If the work during execution displays copyright notices, you must include the copyright notice for the Library among them, as well as a reference directing the user to the copy of this License. Also, you must do one of these things:

- a) Accompany the work with the complete corresponding machine-readable source code for the Library including whatever changes were used in the work (which must be distributed under Sections 1 and 2 above); and, if the work is an executable linked with the Library, with the complete machine-readable "work that uses the Library", as object code and/or source code, so that the user can modify the Library and then relink to produce a modified executable containing the modified Library. (It is understood that the user who changes the contents of definitions files in the Library will not necessarily be able to recompile the application to use the modified definitions.)

- b) Use a suitable shared library mechanism for linking with the Library. A suitable mechanism is one that (1) uses at run time a copy of the library already present on the user's computer system, rather than copying library functions into the executable, and (2) will operate properly with a modified version of the library, if the user installs one, as long as the modified version is interface-compatible with the version that the work was made with.
- c) Accompany the work with a written offer, valid for at least three years, to give the same user the materials specified in Subsection 6a, above, for a charge no more than the cost of performing this distribution.
- d) If distribution of the work is made by offering access to copy from a designated place, offer equivalent access to copy the above specified materials from the same place.
- e) Verify that the user has already received a copy of these materials or that you have already sent this user a copy.

For an executable, the required form of the "work that uses the Library" must include any data and utility programs needed for reproducing the executable from it. However, as a special exception, the materials to be distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

It may happen that this requirement contradicts the license restrictions of other proprietary libraries that do not normally accompany the operating system. Such a contradiction means you cannot use both them and the Library together in an executable that you distribute.

7. You may place library facilities that are a work based on the Library side-by-side in a single library together with other library facilities not covered by this License, and distribute such a combined library, provided that the separate distribution of the work based on the Library and of the other library facilities is otherwise permitted, and provided that you do these two things:

- a) Accompany the combined library with a copy of the same work based on the Library, uncombined with any other library facilities. This must be distributed under the terms of the Sections above.
- b) Give prominent notice with the combined library of the fact that part of it is a work based on the Library, and explaining where to find the accompanying uncombined form of the same work.

8. You may not copy, modify, sublicense, link with, or distribute the Library except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, link with, or distribute the Library is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

9. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Library or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Library (or any work based on the Library), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Library or works based on it.

10. Each time you redistribute the Library (or any work based on the Library), the recipient automatically receives a license from the original licensor to copy, distribute, link with or modify the Library subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties with this License.

11. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Library at all. For example, if a patent license would not permit royalty-free redistribution of the Library by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Library.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply, and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

12. If the distribution and/or use of the Library is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Library under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

13. The Free Software Foundation may publish revised and/or new versions of the Lesser General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Library specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Library does not specify a license version number, you may choose any version ever published by the Free Software Foundation.

14. If you wish to incorporate parts of the Library into other free programs whose distribution conditions are incompatible with these, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

## **NO WARRANTY**

15. BECAUSE THE LIBRARY IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE LIBRARY, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE LIBRARY "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE LIBRARY IS WITH YOU. SHOULD THE LIBRARY PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

16. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE LIBRARY AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE LIBRARY (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE LIBRARY TO OPERATE WITH ANY OTHER SOFTWARE), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

## **END OF TERMS AND CONDITIONS**

### **How to Apply These Terms to Your New Libraries**

If you develop a new library, and you want it to be of the greatest possible use to the public, we recommend making it free software that everyone can redistribute and change. You can do so by permitting redistribution under these terms (or, alternatively, under the terms of the ordinary General Public License).

To apply these terms, attach the following notices to the library. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

one line to give the library's name and an idea of what it does.

Copyright (C) year name of author

This library is free software; you can redistribute it and/or modify it under the terms of the GNU Lesser General Public

License as published by the Free Software Foundation; either version 2.1 of the License, or (at your option) any later version.

This library is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU Lesser General Public License for more details.

You should have received a copy of the GNU Lesser General Public License along with this library; if not, write to the Free Software Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA

Also add information on how to contact you by electronic and paper mail.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a "copyright disclaimer" for the library, if necessary. Here is a sample; alter the names:

Yoyodyne, Inc., hereby disclaims all copyright interest in the library `Frob' (a library for tweaking knobs) written by James Random Hacker.

signature of Ty Coon, 1 April 1990

Ty Coon, President of Vice

That's all there is to it!