

RTView® Enterprise User's Guide

Version 5.3



RTView®

Copyright © 1998-2020. 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®

©2013-2020 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 © 1998-2020 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>

Preface	1
About This Guide	1
Audience	1
Document Conventions	1
Additional Resources	2
Release Notes	2
Support Knowledge Base	2
SL Documentation	2
Contacting SL.....	2
Internet	2
Technical Support.....	2
Chapter 1 - Introduction to RTView Enterprise	3
Displays	4
RTViewCentral	4
RTView DataServers.....	5
Solution Packages	6
System Requirements	6
Chapter 2 - Using the Monitor	7
Login to RTView Enterprise	7
Services Tab.....	8
Service Diagrams.....	8
Defining Flow and Display Names	9
Tips.....	10
Defining Diagram Defaults	10
Viewing Service Diagrams.....	10
Layout Options	11
Components Tab.....	12
Common/Alerts	12
Alerts History Table	12
Alerts Table by Component	13
Alert Detail for Component	14
Alert Configuration for Component.....	16
Common/System	18
System Overview	18
RTView Data Servers	19
Data Server Summary	20
CI Type Definitions	21
CI Stats	22
Cache Table.....	23
Alerts Tab	25
Alerts Table	25

Admin Tab	26
Alert Administration	27
Alert Overrides Administration.....	28
CMDB Administration.....	30
CMDB Administration	31
Add CIs to the CMDB	32
Delete a CI	32
Edit CI Properties	33
Create, Delete, Rename or Merge Owners, Areas, Groups or Services	33
Move a Service, Group or Area	33
Custom Tab.....	36
Fundamental Structure of Displays	36
Heatmaps	37
Mouse-over	37
Tables	38
Multiple Column Sorting	38
Column Visibility.....	39
Column Filtering.....	39
Column Locking.....	41
Column Reordering	41
Saving Settings.....	41
Row Paging	42
Trend Graphs	42
Mouse-over	43
Log Scale	43
Time Range	43
Chapter 3 - RTView DataServer for IBM	45
IBM DB2.....	45
IBM DB2 Overview Display.....	45
DB2 Instances View	47
DB2 Instances Table	47
DB2 Instances Heatmap.....	47
Instance Summary	48
DB2 Members View	51
DB2 Members Table.....	51
DB2 Member Summary	53
DB2 Databases View	55
DB2 Databases Table	55
Single DB2 Database Summary.....	58
Partition Table	61
Partition Heatmap	63
IBM MQ.....	66
IBM MQ Overview Display.....	67

IBM MQ Brokers View	68
IBM MQ Broker Connections Table	69
IBM MQ Brokers Heatmap.....	70
IBM MQ Broker Summary	72
IBM MQ Channels View	75
IBM MQ Channels Table.....	76
IBM MQ Channels Heatmap.....	79
IBM MQ Channel Summary	82
IBM MQ Queues View	85
IBM MQ Queue Managers Table	85
IBM MQ Queues Table	87
IBM MQ Queues Heatmap.....	89
IBM MQ Queue Summary	92
IBM MQ Topics View	95
IBM MQ Topics Table.....	95
IBM MQ Topics Heatmap.....	97
IBM Topic Summary	99
IBM MQ Topic Configs View.....	103
IBM MQ Topic Configs Table.....	103
IBM MQ Topic Config Summary	105
IBM MQ Topic Publishers View.....	107
IBM MQ Topic Publishers Table	107
IBM MQ Topic Publishers Heatmap	109
IBM MQ Topic Publisher Summary	111
IBM MQ Topic Subscribers View.....	114
IBM MQ Topic Subscribers Table	114
IBM MQ Topic Subscribers Heatmap	116
IBM MQ Topic Subscriber Summary	118
IBM MQ Topic Publisher Totals View	122
IBM MQ Topic Publisher Totals Table	122
IBM MQ Topic Publisher Totals Heatmap	124
IBM MQ Topic Publisher Totals Summary	125
IBM MQ Topic Subscriber Totals View	128
IBM MQ Topic Subscriber Totals Table	128
IBM MQ Topic Subscriber Totals Heatmap	130
IBM Topic Subscriber Totals Summary.....	131
IBM MQ Subscriptions View	134
IBM MQ Subscriptions Table.....	134
IBM MQ Subscriptions Heatmap.....	136
IBM MQ Subscription Summary	138
IBM MQ Subscription Configs View	141
IBM MQ Subscription Configs Table	142
IBM MQ Subscription Config Summary.....	144
IBM WebSphere.....	147
WebSphere Overview.....	147
WebSphere Servers Heatmap.....	149
WebSphere Server Summary	149

WebSphere Apps Table	150
WebSphere Apps Heatmap	151
WebSphere App Summary	151
Chapter 4 - RTView DataServer for Infrastructure	153
Amazon Web Services	153
Amazon EC2 Overview	154
Amazon EC2 Instance Heatmap.....	155
Amazon EC2 Instance Table.....	157
Amazon EC2 Instance Summary	159
Docker	161
Docker Overview	161
Docker Engines View	162
Docker Engines Table.....	163
Docker Engines Heatmap.....	164
Docker Single Engine Summary.....	167
Docker Containers View	170
Docker Containers Table.....	170
Docker Containers Heatmap.....	172
Docker Single Container Summary.....	174
JBoss	178
JBoss Overview	178
JBoss Servers	178
JBoss Servers Table.....	178
JBoss Servers Heatmap.....	179
JBoss Single Server Summary	179
JBoss Applications.....	179
JBoss Apps Table.....	179
JBoss Apps Heatmap.....	180
JBoss App Totals	180
MongoDB.....	181
MongoDB Overview	181
Mongo Instances View	182
MongoDB Instances Table.....	183
MongoDB Instances Heatmap.....	183
MongoDB Instance Summary	186
Mongo Databases View	187
MongoDB Databases Table.....	187
MongoDB Databases Heatmap.....	188
MongoDB Single Database Summary.....	190
Mongo Collections View	191
MongoDB Collections Table	191
MongoDB Collections Heatmap	192
MongoDB Single Collection Summary	194
MySQL Database	195

MySQL Overview Display	195
All MySQL Instances View	196
Instances Heatmap.....	196
All MySQL Instances	199
Single MySQL Instance View.....	199
Single MySQL Instance Summary	200
Instance Properties.....	201
Instance Operations.....	201
Instance User Tables.....	202
MS SQL	203
MS SQL Overview	204
MS Servers View.....	205
MS SQL Servers Table.....	206
MS SQL Servers Heatmap.....	208
MS SQL Server View.....	209
MS SQL Server Summary	210
MS SQL Server Performance	215
MS SQL Server Waits Table.....	218
MS SQL Server DB Table Sizes	220
Node.js	221
Node.js Overview.....	222
Node/Master View.....	223
Node/Master Info Table	224
Node Master Summary.....	226
Node Requests View.....	229
Node Requests Table	229
Node Request Summary.....	231
Node Processes View	233
Node Processes Table	234
Node Processes Heatmap	236
Node Process Summary.....	238
RTView Host Agent.....	242
Hosts Overview	242
Hosts Heatmap	243
Hosts Table	245
Host Summary.....	246
Host Processes.....	246
Host Network Interfaces.....	247
Host Storages	248
VMware vCenter	249
VMware Overview Display	250
VMware Hosts View	251
VMware Hosts Table	252
VMware Host Summary	255
VMware Machines View.....	259
VMware Virtual Machines Table	259

VMware Virtual Machine Summary	262
Chapter 5 - RTView DataServer for Kafka	267
Kafka Overview	268
Kafka Clusters View	269
Kafka Clusters Table	270
Kafka Single Cluster Performance	271
Kafka Brokers View	273
Kafka Brokers Table	273
Kafka Brokers Heatmap	275
Kafka Single Broker Summary	279
Kafka Single Broker JVM Runtime Summary	283
Kafka Single Broker Topics Summary	286
Kafka Single Broker Topic Lag Summary	288
Kafka Zookeepers View	289
Kafka Zookeepers Table	290
Kafka Zookeepers Heatmap	292
Kafka Single Zookeeper Summary	295
Kafka Single Zookeeper JVM Runtime Summary	300
Kafka Topics View	304
Kafka Topics Table	304
Kafka Topics Heatmap	306
Kafka Single Topic Summary	309
Kafka Single Topic Partition Summary	313
Kafka Topics Activity by Cluster	314
Kafka Brokers Activity by Topic	316
Kafka Producers View	318
Kafka Producers Table	319
Kafka Producers Heatmap	322
Kafka Single Producer Summary	326
Kafka Single Producer JVM Runtime Summary	330
Kafka Consumers View	333
Kafka Consumers Table	333
Kafka Consumers Heatmap	336
Kafka Single Consumer Summary	339
Kafka Single Consumer JVM Runtime Summary	343
Kafka Single Consumer Lag Summary	347
Chapter 6 - RTView DataServer for Oracle	349
Oracle Coherence	349
Oracle Database Overview Display	349
Oracle Databases View	351
Oracle Databases Table	351
Databases Heatmap	352
Databases Summary	355
Oracle Instances View	356

Oracle Instances Table	356
Instances Heatmap	357
Instance Summary	359
Oracle Database	360
Oracle Database Overview Display	360
Oracle Databases View	361
Oracle Databases Table	361
Databases Heatmap	363
Databases Summary	365
Oracle Instances View	366
Oracle Instances Table	366
Instances Heatmap	367
Instance Summary	369
Connector for Oracle Enterprise Manager	370
Oracle WebLogic	370
WebLogic Overview Display	370
WebLogic Servers View	372
WebLogic Servers Table	372
WebLogic Servers Heatmap	374
WebLogic Clusters Table	377
Single WebLogic Server View	378
WebLogic Server Summary	379
WebLogic JVM Summary	383
WebLogic Server Detail Tables	386
WebLogic JDBC Summary	390
WebLogic Server Work Manager Table	395
WebLogic JMS Persistent Stores Detail Tables	396
WebLogic Applications View	398
WebLogic Cluster Applications Table	399
WebLogic Clustered Application Summary	400
WebLogic Server Applications Heatmap	403
WebLogic Server Applications Summary	406
WebLogic Application Summary	409
WebLogic Application Metric Trends	414
WebLogic Application Components Heatmap	417
WebLogic Application Component Summary	420
Chapter 7 - RTView DataServer for Solace	423
Brokers	424
Brokers Overview	424
Brokers Heatmap	426
Brokers Table	428
Broker Summary	435
Broker Sensors	437
Broker Provisioning	438

Broker Interface.....	440
Brokers Message Spool	442
CSPF Neighbors	444
Neighbors Table	444
CSPF Neighbors Diagram.....	446
Neighbors Summary	448
VPNs.....	450
VPNs Heatmap.....	450
VPNs Table.....	454
VPNs Summary	457
Clients.....	460
Clients Table.....	460
Client Summary	465
Bridges.....	468
Bridges Table.....	468
Bridges Diagram	471
Bridge Summary	472
Endpoints	475
Endpoints Table	475
Endpoint Summary	477
Capacity	480
Broker Capacity Table.....	480
Broker Capacity - Summary.....	482
Broker Capacity Trends	483
Syslog Events.....	484
Syslog Events Table.....	485
Syslog Event Summary	486
.....	489
Chapter 8 - RTView DataServer for TIBCO.....	491
TIBCO ActiveMatrix	491
TIBCO ActiveSpaces.....	492
Grids Views.....	494
TIBCO ActiveSpaces Grids Table	494
TIBCO ActiveSpaces Grids Heatmap	495
TIBCO ActiveSpaces Grid Summary	497
TIBCO ActiveSpaces Realm Server	499
Nodes Views	501
TIBCO ActiveSpaces Nodes Table.....	502
TIBCO ActiveSpaces Nodes Heatmap.....	505
TIBCO ActiveSpaces Node Summary	507
Proxies Views	510
TIBCO ActiveSpaces Proxies Table	510
TIBCO ActiveSpaces Proxies Heatmap	513
TIBCO ActiveSpaces Proxy Summary	515
Keepers Views	519

TIBCO ActiveSpaces StateKeepers Table.....	519
TIBCO ActiveSpaces StateKeepers Heatmap.....	521
TIBCO ActiveSpaces Keeper Summary	523
TIBCO ActiveSpaces (2.x)	525
TIBCO ActiveSpaces 2 Overview	525
MetaSpaces	527
TIBCO ActiveSpaces 2 Metaspaces Table	527
TIBCO ActiveSpaces 2 MetaSpace Summary	529
Spaces	533
TIBCO ActiveSpaces 2 Spaces Table.....	533
TIBCO ActiveSpaces 2 Spaces Heatmap.....	535
TIBCO ActiveSpaces 2 Space Summary	537
Members	541
TIBCO ActiveSpaces 2 All Members Table.....	541
TIBCO ActiveSpaces 2 Members Heatmap.....	544
TIBCO ActiveSpaces 2 Member Summary	547
TIBCO Adapters.....	551
TIBCO Adapters Overview.....	551
All Adapters View.....	552
TIBCO Adapters Table	553
TIBCO Adapters Heatmap.....	554
TIBCO Adapter Summary	557
TIBCO BusinessEvents.....	561
TIBCO BusinessEvents Overview.....	561
BE Clusters	563
TIBCO BE Clusters Table.....	563
TIBCO BE Clusters Heatmap	565
TIBCO BE Cluster Summary.....	568
BE Nodes.....	571
TIBCO BE Cluster Nodes Table	571
TIBCO BE Cluster Nodes Heatmap	573
TIBCO BE Inference Node Summary.....	575
TIBCO BE Storage Node Summary.....	580
BE Events	585
TIBCO BE Events Table	585
TIBCO BE Agent Event Summary.....	588
TIBCO BE Event Cache Hits Table	592
TIBCO BE Event Hit Summary	594
BE Concepts.....	598
TIBCO BE Concepts Table	598
TIBCO BE Concept Hit Summary.....	600
TIBCO BE Channel Status Table.....	603
TIBCO BE Inference Agents Table	605
TIBCO BE RTC Txn Manager Reports	608
TIBCO BusinessWorks.....	610
TIBCO BusinessWorks Overview Display.....	611

BW Applications.....	613
TIBCO BusinessWorks Applications Table	613
TIBCO BusinessWorks Applications Heatmap	615
TIBCO BusinessWorks Application Summary.....	618
BW Containers.....	621
TIBCO BusinessWorks Containers Table	621
TIBCO BusinessWorks Containers Heatmap	624
TIBCO BusinessWorks Container Summary	627
BW Application Nodes.....	630
TIBCO BusinessWorks Application Nodes Table	630
TIBCO BusinessWorks Application Nodes Heatmap	633
TIBCO BusinessWorks Application Node Summary	635
BW Application Slices	639
TIBCO BusinessWorks Application Slices Table.....	639
TIBCO BusinessWorks Application Slices Heatmap.....	642
TIBCO BusinessWorks Application Slice Summary	644
BW Processes	649
TIBCO BusinessWorks Processes Table	650
TIBCO BusinessWorks Processes Heatmap	653
TIBCO BusinessWorks Process Summary	656
BW Activities.....	660
TIBCO BusinessWorks Activities Table	660
TIBCO BusinessWorks Activities Heatmap	663
TIBCO BusinessWorks Activity Summary	666
TIBCO BusinessWorks 5 Monitor	670
TIBCO BusinessWorks 5 Overview Display	671
BW5 Servers	673
TIBCO BusinessWorks 5 Servers Table	673
TIBCO BusinessWorks 5 Servers Heatmap	674
TIBCO BusinessWorks 5 Server Summary.....	677
BW5 Engines.....	679
TIBCO BusinessWorks 5 Engines Table	680
TIBCO BusinessWorks 5 Engines Heatmap	682
TIBCO BusinessWorks 5 Engine Summary.....	686
BW5 Processes	690
TIBCO BusinessWorks 5 Processes Table	690
TIBCO BusinessWorks 5 Processes Heatmap	693
TIBCO BusinessWorks 5 Process Summary.....	697
BW5 Activities	700
TIBCO BusinessWorks 5 Activities Table.....	700
TIBCO BusinessWorks 5 Activities Heatmap	703
TIBCO BusinessWorks 5 Activity Summary.....	706
TIBCO Enterprise Message Service	710
TIBCO EMS Overview	710
EMS Servers	712
TIBCO EMS Servers Table.....	712

TIBCO EMS Servers Heatmap.....	715
TIBCO EMS Server Summary	719
TIBCO EMS Server Trends	724
TIBCO EMS Server Destinations.....	727
TIBCO EMS Server Connections.....	731
TIBCO EMS Bridges, Users, Ports.....	733
EMS Topics.....	735
TIBCO EMS Topics Table.....	736
TIBCO EMS Topics for Server Summary	740
TIBCO EMS Topics Heatmap.....	743
TIBCO EMS Topic Summary	748
TIBCO EMS Topic Detail by Server	751
EMS Queues.....	753
TIBCO EMS Queues Table	754
TIBCO EMS Queues for Server Summary	757
TIBCO EMS Queues Heatmap	760
TIBCO EMS Queue Summary	764
TIBCO EMS Queue Detail By Server	767
EMS Routes.....	769
TIBCO EMS Routes	769
TIBCO EMS Route.....	771
EMS Producers	774
TIBCO EMS Producers	774
TIBCO EMS Producer.....	777
EMS Consumers.....	779
TIBCO EMS Consumers	779
TIBCO EMS Consumer	783
EMS Durables.....	787
TIBCO EMS Durables.....	787
TIBCO EMS Durable	789
TIBCO FTL	792
TIBCO FTL Overview	792
FTL Servers.....	793
FTL Servers Heatmap.....	794
FTL Servers Table.....	796
FTL Server Summary	797
FTL Clients.....	798
All Clients Heatmap	798
All Clients Table	800
Single Client Summary	801
Chapter 9 - RTView Manager.....	805
Login to RTView Manager	805
Displays	807
Tomcat Displays.....	807
Tomcat Overview	808

Tomcat Servers Heatmap	809
Single Tomcat Server	810
.....	810
All Tomcat Apps	811
Single Tomcat App	812
JVM Processes Displays	814
JVM Overview	814
JVMs Table	815
JVMs Heatmap	816
JVM Summary	818
JVM System Properties	819
JVM GC Trends	820
RTView Servers Displays	821
Data Servers	822
Data Server Summary	823
Historian Servers	823
'Drilldowns' Displays	824
Alerts History Table	824
Alerts Table by Component	825
Alert Detail for Component	827
Alerts Displays	828
Alerts Table	828
Admin Displays	830
Alert Administration	830
Alert Overrides Admin	832
Cache Table	834
Alerts for RTView Manager	835
Configure High Availability	837
Appendix A - Alert Definitions	839
Amazon Web Services	839
Apache Kafka	841
Docker	845
IBM MQ	846
Microsoft SQL Server	849
MongoDB	850
MySQL Database	852
Node.js	853
Oracle Coherence	854
Oracle Database	861
Oracle WebLogic	863

RTView Host Agent.....	867
RTView Manager and RTView Rules	869
Solace.....	870
TIBCO ActiveMatrix BusinessWorks.....	878
TIBCO ActiveSpaces	883
TIBCO ActiveSpaces (2.x)	886
TIBCO Adapters.....	888
TIBCO BusinessEvents.....	889
TIBCO Enterprise Message Service	891
TIBCO FTL	896
UX.....	898
VMware vCenter	899
Appendix B - Limitations	903
iPad Safari Limitations	903
TIBCO ActiveMatrix BusinessWorks.....	905
Servers.....	905
Business Works 5.7.1 Engine Status	905
BWSE Components	906
Appendix C - Third Party Notice Requirements	907
RTView Enterprise.....	907
RTView Core	913

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.
boldface	Within text, directory paths, file names, commands and GUI controls appear in bold typeface.
Courier	Code examples appear in Courier font: <code>amnesiac > enable</code> <code>amnesiac # configure terminal</code>
< >	Values that you specify appear in angle brackets: <code>interface <ipaddress></code>

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 the most current version of SL documentation, visit the SL [Documentation](#) Web site. For a complete list of SL documentation, visit the SL Support Web site located at <http://www.sl.com/support/>.

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

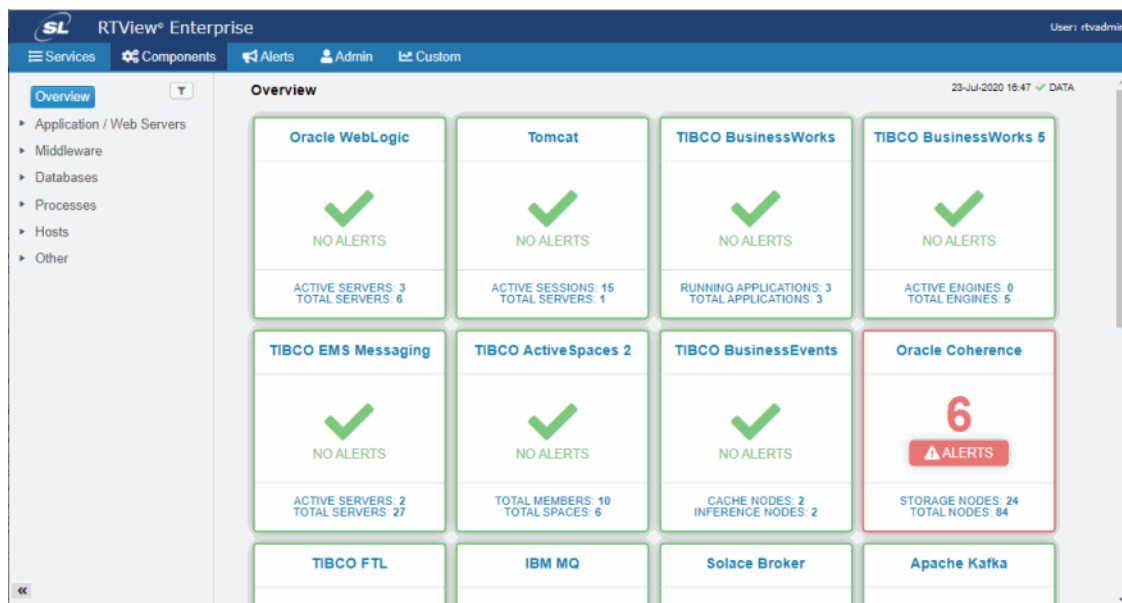
If you have problems installing, using, or replacing SL products, contact SL Support or your channel partner who provides support. To contact SL Support, open a trouble ticket by calling 415 927 8400 in the United States and Canada or +1 415 927 8400 outside the United States.

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.

This chapter describes Enterprise Monitor displays, GUI and behavior. The following figure illustrates the RTView® Enterprise Components tab Overview page.



For details about setting up RTView Enterprise, RTViewCentral, RTView DataServers and solution packages, see the *RTView Enterprise Configuration Guide*, available on our [Documentation](#) website.

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 RTVRules.


RTViewCentral displays come with RTView Enterprise and reside on RTViewCentral.

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 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, JVMs, Tomcat 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.

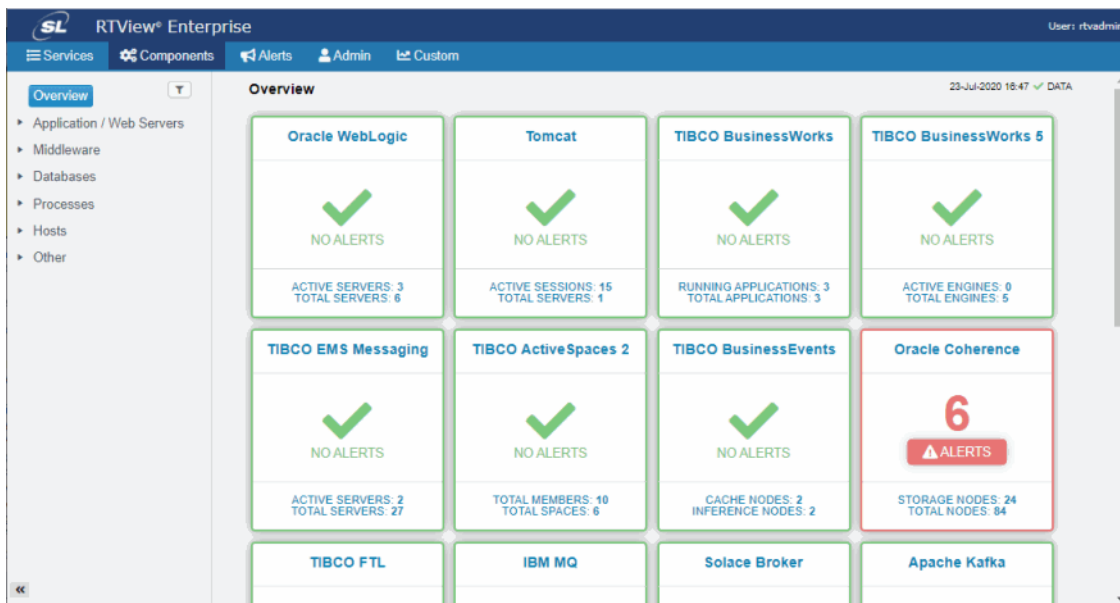
CHAPTER 2 Using the Monitor

Welcome to RTView Enterprise. This section describes how to access the Monitor and the many displays that come with RTViewCentral.

Login to RTView Enterprise

Browse to: <http://localhost:11070/rtview-central> (login as rtvadmin/rtvadmin or rtvuser/rtvuser).

The **Overview** page (shown below), located in the **Components** tab, opens by default. Notice the navigation tabs in the title bar--the “[Services Tab](#)”, “[Components Tab](#)”, “[Alerts Tab](#)”, “[Admin Tab](#)” and the “[Custom Tab](#)”--this section describes each of them.



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

This section contains:

- [“Services Tab”](#)
- [“Components Tab”](#)
- [“Alerts Tab”](#)
- [“Admin Tab”](#)
- [“Custom Tab”](#)
- [“Fundamental Structure of Displays”](#)
- [“Heatmaps”](#)
- [“Tables”](#)
- [“Trend Graphs”](#)

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

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

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

Services Tab

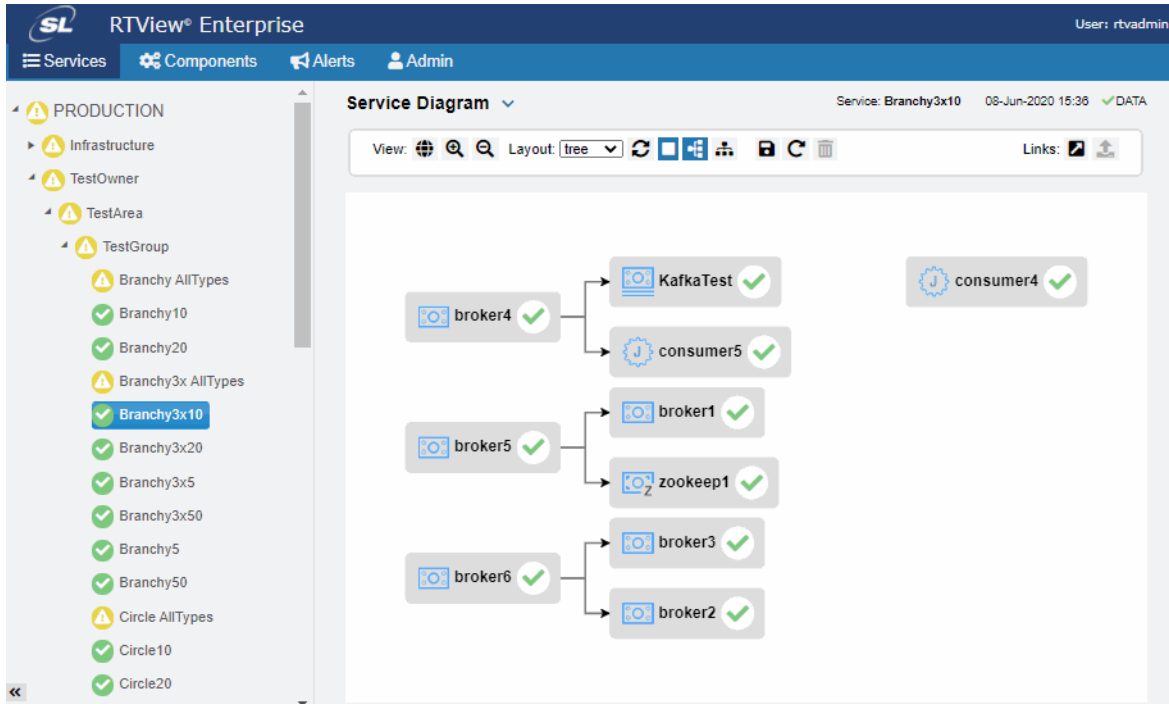
Provides a tree structure view of your defined CMDB with four levels of hierarchy: **Owner>Area>Group>Service**. The tree is configurable and should be set up to represent how a support person likes to conceptually think of the vast number of items that are being monitored.

The Services tab features [“Service Diagrams”](#).

Service Diagrams

Service flow diagrams show all components within a selected service. Each node in the diagram represents a single component. The label for each node is either the **Display Name** (if specified) or a shortened version of the **CI Name**. The icon (above or to the left of the label) allows you to quickly identify the CI Type and the alert indicator (below or to the right of the label) shows the maximum alert level for the component.

To navigate to the service flow diagram, go to the **Services** tab and click on a single service in the service tree. The **CI Stats Table** opens. Click the arrow to next to the **CI Stats Table** title and select **Service Diagram**. You see one node for each row in the **CI Stats Table** in the selected service. Until you have defined relationships between the components, all components are laid out in a grid.





You can mouse over the node to see the **CI Name**, **CI Type**, **Quality**, **Alert Level** and **Alert Count**. If the CI Type supports key metrics, the mouse over also shows the names and values of all of the top level key metrics. Double-click on the node to drill down to the associated component summary display.



Defining Flow and Display Names

You can define the service flow after you configure the service model. The **Infrastructure** service model does not support relationships between the components. To define the flow of the components within a service, you must be logged in as `rtvadmin`.

To define flow and display names:

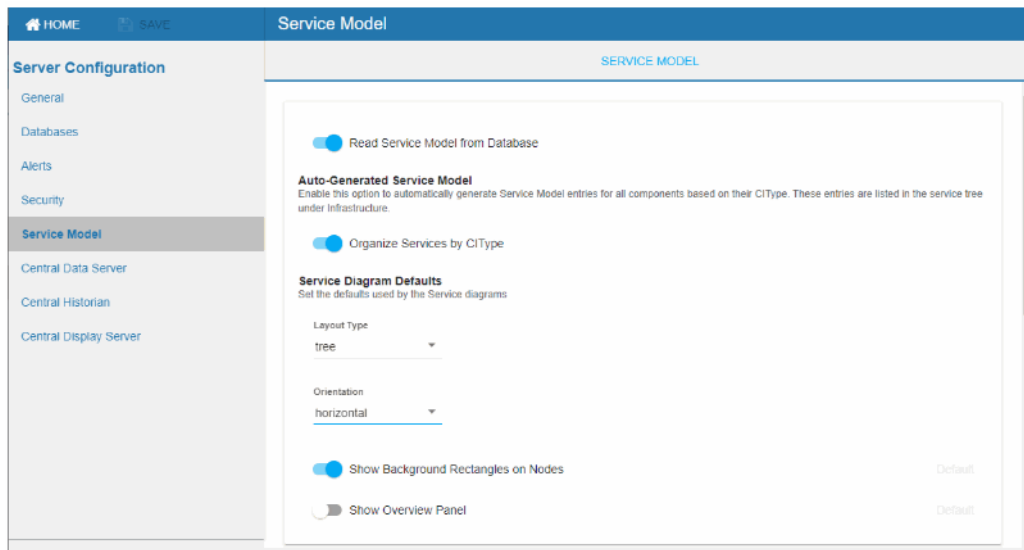
1. Navigate to the flow diagram of the service you want to edit.
2. Click  then select a node and drag it to another node. This draws a directional link from the first node to the second node. To delete an existing link, right-click on the link and then click **Delete** in the popup menu. After you create or delete one or more links, click  to save the changes to the service model.

Tips

- with the link button  disabled, you can click and drag nodes to rearrange them in order to make it easier to draw links between them
- at any point while you are editing, you can click  to re-apply the selected layout options to your diagram
- Component names are limited to 15 characters or less for use in the labels. If the component name is longer than 15 characters, contains a semi-colon and the section after the last semi-colon is less than 15 characters, that is used. Otherwise, the first and last 6 characters are used separated by a ~. To specify an alternative name to use, go the **CMDB Admin** page on the Admin, select the component within the service, click **Edit**, enter a **Display Name** and save your changes.

Defining Diagram Defaults

By default, diagrams use a horizontal tree layout with backgrounds rectangles on all nodes and the overview panel is off. You can change these defaults on the Service Model tab of the Configuration Application. The layout defaults are used for all diagrams and for all users unless a user has locally saved a layout for a diagram.



Viewing Service Diagrams

To navigate to the service diagram, go to the **Services** tab and click a single service in the service tree. This opens the **CI Stats Table**. Click the arrow next to the **CI Stats Table** title and select **Service Diagram**. Until relationships between the components have been defined in the Service Model, all components are laid out in a grid.

The toolbar above the diagram allows you to interact with the diagram:





View Options

 - Toggle the visibility of the overview panel.

 - Zoom In


 - Zoom Out

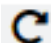
 - Apply the selected layout options. This is useful to restore the layout after clicking and dragging the nodes or while adding links.

 - Toggle the visibility of the icon background rectangles.

 - Use a horizontal flow (applies to tree and graph layout only).

 - Use a vertical flow (applies to tree and graph layout only).


 - Save the layout options including the positions of any manually moved nodes to the local browser storage. This will be used instead of the default layout whenever you view this diagram in the browser where it was saved.

 - Restore the layout options in the toolbar to the locally saved layout if there is one, otherwise to the default layout.

 - Delete the locally saved layout options.

Layout Options

You can select a layout type to apply to the diagram. You can also click and drag the component nodes to lay them out manually.



The **tree** and **graph** layouts are useful on diagrams where the flow (links) between nodes go mostly in one direction, from one or more root nodes to child nodes, then to grandchild nodes, and so forth. The graph layout attempts to keep the nodes arranged neatly in rows or columns, the tree layout places the nodes more closely together.

The **net** layout can be useful for diagrams where the flow between nodes goes in multiple directions.

The **circle** layout can be useful for diagrams where the flow is mostly in one direction but where there are also links from the deepest child nodes back to top level nodes.

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.

The Components tab also has the following Views and displays in the navigation tree:

- **“Common/Alerts”**: 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”**: These displays show details about your RTView Enterprise system.

Common/Alerts

This View includes the following displays:

- **“Alerts History Table”**: Track history of any alert that has occurred in your RTView Enterprise system.
- **“Alerts Table by Component”**: Track alerts associated with CIs shown in a display.
- **“Alert Detail for Component”**: Investigate an alert instance and its history.
- **“Alert Configuration for Component”**: Refine alert threshold settings.

Alerts History Table

Use this display to track the history of alerts, including cleared alerts in your monitoring system. There is one row in the table for each update to each alert.

Choose a Data Server from the drop down to filter alerts shown in the table. The **Alerts History Table** only shows alerts associated with the selected Data Server.

Select **Expand Alert Index** to separate each column in the **Alert Index** into different lines of text. When unselected, the **Alert Index** remains as a single line, with all index parts separated by semicolon (;).

Select **History Alerts** to show all historical alerts. When unselected, only current alerts are shown in the table.

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**.

Alerts History Table 07-Oct-2019 09:25 DATA

Data Server: RTV-DATA-TIBCO Expand Alert Index: History Alerts:

Data Server URL: https://rtvdemos.sl.com/tibmon_rtviewquery/ 15 minutes

Alert Level	Ack	Cleared	Alert Name	Alert Index	Alert Text	Owner	Id	Source	Row Update Time
			JvmMemoryUsed-High	win44-newbe4cache	High Alert Limit exceeded		144843		2019-Oct-07 09:11:41
		<input checked="" type="checkbox"/>	JvmMemoryUsed-High	win44-newbe4cache	High Alert Limit exceeded		144843		2019-Oct-07 09:11:50
			BwServerCpuUsed-High	sls4-64(slmon)	High Alert Limit exceeded		144911		2019-Oct-07 09:18:4
		<input checked="" type="checkbox"/>	BwServerCpuUsed-High	sls4-64(slmon)	High Alert Limit exceeded		144911		2019-Oct-07 09:19:0
			BwProcessAvgElapsedTimeHk	sls4-64(slmon)-domains	High Alert Limit exceeded		144855		2019-Oct-07 09:10:5
			BwProcessAvgElapsedTimeHk	sls4-64(slmon)-domains	High Alert Limit exceeded		144854		2019-Oct-07 09:10:5
			BwProcessAvgElapsedTimeHk	sls4-64(slmon)-domains	High Alert Limit exceeded		144853		2019-Oct-07 09:10:5
			BwProcessAvgElapsedTimeHk	sls4-64(slmon)-domains	High Alert Limit exceeded		144852		2019-Oct-07 09:10:5
			BwProcessAvgElapsedTimeHk	sls4-64(slmon)-domains	High Alert Limit exceeded		144851		2019-Oct-07 09:10:5
			BwProcessAvgElapsedTimeHk	sls4-64(slmon)-domains	High Alert Limit exceeded		144856		2019-Oct-07 09:11:5
			BwProcessAvgElapsedTimeHk	slhpux11(slmon)-domain	High Alert Limit exceeded		144865		2019-Oct-07 09:12:1
			BwProcessAvgElapsedTimeHk	slhpux11(slmon)-domain	High Alert Limit exceeded		144864		2019-Oct-07 09:12:1
			BwProcessAvgElapsedTimeHk	slhpux11(slmon)-domain	High Alert Limit exceeded		144863		2019-Oct-07 09:12:1
			BwProcessAvgElapsedTimeHk	slhpux11(slmon)-domain	High Alert Limit exceeded		144862		2019-Oct-07 09:12:1
			BwProcessAvgElapsedTimeHk	slhpux11(slmon)-domain	High Alert Limit exceeded		144861		2019-Oct-07 09:12:1
			BwProcessAvgElapsedTimeHk	slhpux11(slmon)-domain	High Alert Limit exceeded		144860		2019-Oct-07 09:12:1
			BwProcessAvgElapsedTimeHk	slhpux11(slmon)-domain	High Alert Limit exceeded		144859		2019-Oct-07 09:12:1
			BwProcessAvgElapsedTimeHk	sls4-64(slmon)-domains	High Alert Limit exceeded		144872		2019-Oct-07 09:12:1
			BwProcessAvgElapsedTimeHk	sls4-64(slmon)-domains	High Alert Limit exceeded		144871		2019-Oct-07 09:12:1
			BwProcessAvgElapsedTimeHk	sls4-64(slmon)-domains	High Alert Limit exceeded		144870		2019-Oct-07 09:12:1
			BwProcessAvgElapsedTimeHk	sls4-64(slmon)-domains	High Alert Limit exceeded		144869		2019-Oct-07 09:12:1
			BwProcessAvgElapsedTimeHk	sls4-64(slmon)-domains	High Alert Limit exceeded		144868		2019-Oct-07 09:12:1

Alerts Table by Component

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.

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: False ACK: False

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_lqa)	High V
2018-Oct-01 06:20:10				⚠️	HostCpuPercentHigh	SL-DEMO:SLHOST17(sl_lamx)	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_lamx)	High A
2018-Oct-01 06:18:36				🔔	HostMemoryUsedHigh	SL-DEMO:SLHOST17(sl_lamx)	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 10:33:51				⚠️	HostVirtualMemoryUsedH	SL-DEMO:SLHOST16(sl_lqa)	High A
2018-May-01 03:45:4				⚠️	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

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	🔔				

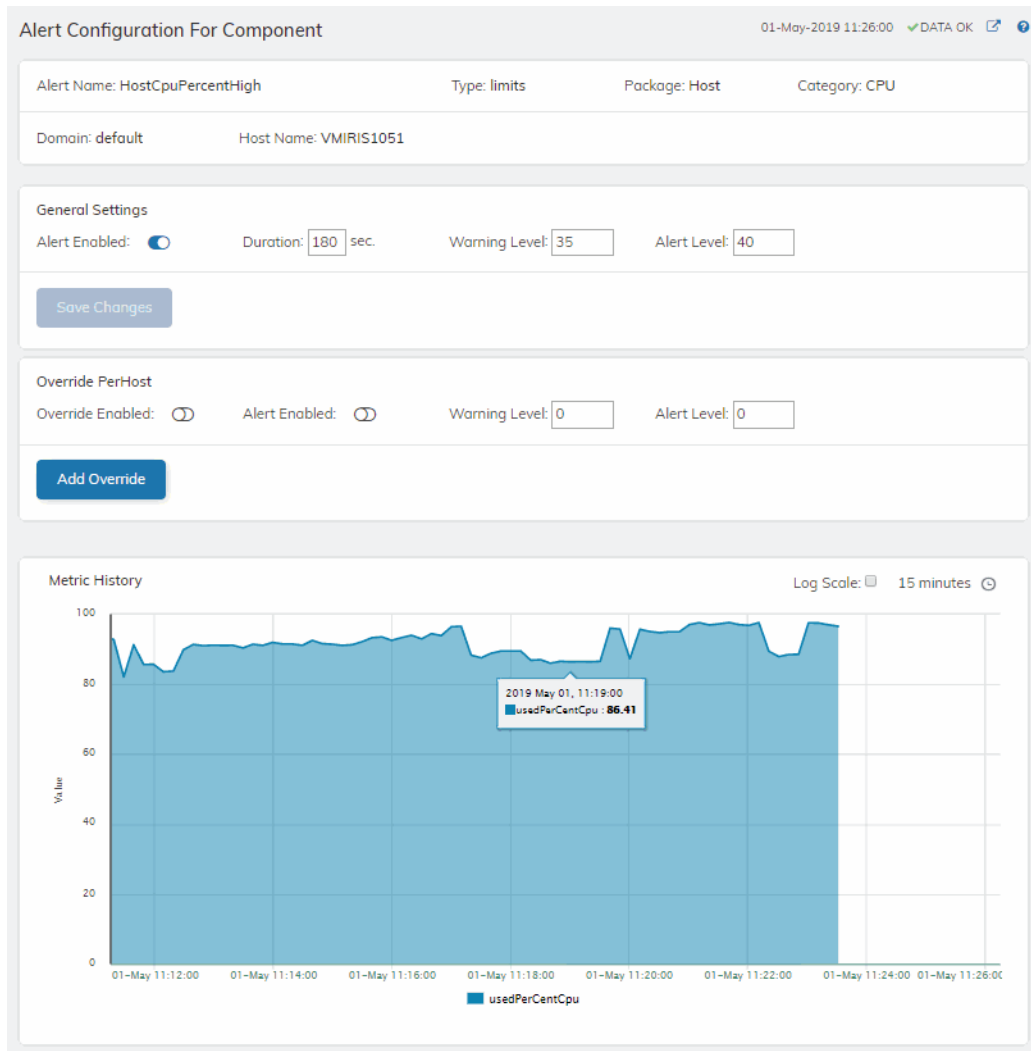
Alert Configuration for Component

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

RTView Enterprise includes the following displays which provide details about your RTView system setup:

- ["System Overview"](#)
- ["RTView Data Servers"](#)
- ["Data Server Summary"](#)
- ["CI Type Definitions"](#)
- ["CI Stats"](#)
- ["Cache Table"](#)

System Overview

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.
- **CI Type 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_rtv	BOOMI-ATOM JVM
RTV-DATA-IBM	✓	0	85	https://rtvdemos.sl.com/ibmmon_rtvque	DB2 MQ-BROKER MQ-CHANNEL MQ
RTV-DATA-INFRA	✓	0	417	https://rtvdemos.sl.com/inframom_rtvqu	DOCKER-CONTAINER DOCKER-ENG
RTV-DATA-KAFKAMON	✓	185	142	https://rtvdemos.sl.com/kafkamom_rtvqu	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_rtvque	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

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” 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

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/tibmon_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/solmon_rtvddata/	1504	02-May-2019 16:54:59	AP
RTV-DATA-IBM	✓	OK	https://rtvdemos.sl.com/ibmmom_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

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

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 **CI Types** for which we are receiving data, provides the associated **CACHENAME** and the associated **Source** RTView DataServer.

The **Cache Map By CI** table lists all **CIs** 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	slsl4-84(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

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: History Tables:

Data Server URL: 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	8	12506

Cache: RtvAlertStatsByCIAndAlertGroup Table: current

time_stamp	CITYPE	CINAME	ALERTGROUP	MaxSeverity	AlertCount
2019-May-07 14:11:33	JVM	localhost:SOLMON_TOM	None	2	1
2019-May-07 14:11:33	JVM	localhost:EMSMON_TOI	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

Alerts Tab

Provides access to the **Alerts Table** where you can manage alerts by owning them, acknowledging them, and/or suppressing them.

Alerts Table

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), **Clear** to set the **Cleared** flag on the selected alert(s), **Comment** to add a comment to the alert(s) and **CI** to get details about the CI associated with the alert (these buttons are enabled when you click one or more alerts).

You must be logged in as `rtvalertmgr` or `rtvadmin` to perform the **Own**, **Ack**, **Unack**, or **Comment** actions. Otherwise, you get an error dialog.

Time	Ack	Ctr	Sevl	Alert Name	Alert Text	Owner	ID	Source	Comments	CIn
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 exoceed		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 exoceed		1011	Z-SIMDATA-1		defa

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.

These displays enable you to set alert thresholds, and observe how alerts are managed. Displays in this View are:

- [“Alert Administration”](#)
- [“Alert Overrides Administration”](#)
- [“CMDB Administration”](#)

Alert Administration

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.

You can also override the alert duration time per alert index instead of to all indexes. To override the duration for an alert index, select the alert in the **Alert Administration** display, click **Override** and edit the **Alert Delay**. For alert indexes that were overridden in a previous release (before duration override was supported) the override duration is set to **-1**, indicating that this is set to use the top level alert duration.

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.

The screenshot shows the 'Alerts Administration' interface. At the top, there is a package dropdown set to 'All' and a URL 'http://rtvdemos.sl.com/emdemo_central_rtvquery'. Below this is a table with the following columns: Alert Name, Alert Enabled, Alert Delay, Warning Level, Alert Level, and Override Count. The table lists various alerts, with 'JvmThreadCountHigh' selected. Below the table is a settings panel for the selected alert, 'HostSwapUsedHigh', with fields for Delay (30), Warning Level (75), and Alert Level (90). There are buttons for 'Save Settings', 'Original Defaults', and 'Override Settings'. At the bottom, it shows 'Alert Selected: HostSwapUsedHigh' and its description: 'The percentage of swap space used is above the limits defined for that Host'.

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: 30 Warning Level: 75 Alert Level: 90

Buttons: Save Settings, Original Defaults, Override Settings

Alert Selected: HostSwapUsedHigh Description: The percentage of swap space used is above the limits defined for that Host

See [“Alert Overrides Administration”](#) for additional details.

Alert Overrides Administration

Administrators use this display to override the alert settings defined in the **Alert Administration** display. To access this display, select an alert in the **Alert Administration** display and choose **Override Settings**.

The table lists all the possible overrides that can be defined for the alert you selected from the **Alert Administration** display. Each row in the table represents a different resource or group of resources that can be overridden. When the four last columns are blank, that means the resource has not been overridden, and the default settings for the alert apply. Otherwise, columns describe whether the alert is enabled, if the override itself is enabled, the overridden alert thresholds and the overridden duration for each row.

Use the **Override Type** drop-down menu to switch the list to a specific type of override (the options for this menu vary according to the alert type), and use the **Display** drop-down menu to list **All** resources, **Overridden** resources or **Free** resources.

You can also enter a pattern or regular expression in the **Search** string to limit the list.

The **RegEx** checkbox indicates whether the text you entered is treated as a search pattern or as a regular expression. Multiple rows can be selected to create/edit/remove many overrides simultaneously.

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 overrides:

Select an **Override Type** from the drop-down menu (depending on the alert, there might be only one type) and then select one or more rows from the table. Under **Settings for selected index** (in the lower portion of the screen), modify settings for the **Override Enabled**, **Alert Enabled**, **Alert Delay**, **Warning Level** and/or **Alarm Level**, then click **Add Override**. The table updates with your new settings.

To alter overrides:

To alter existing overrides with new settings, select them from the table, set all properties under **Settings for selected index** as desired, then click **Save Settings**. To clear existing overrides, select one or more rows, then click **Remove Override**.

Note: You can override alert and warning levels without overriding duration by setting it to **-1**.

For alert indexes that were overridden in a previous release (before duration override was supported) the override duration is set to **-1**, indicating that this is set to use the top level alert duration.

Alert Overrides Administration Data Server: TB-DataServerInfra 24-Jun-2020 14:43:58 DATA OK

Alert: **AcwInstanceDiskReadOpsHigh** Override Type: **PerInstance** Display: **All**

Search: RegEx:

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:

CMDB Administration

This display allows you to modify your Service Data Model.

- ["CMDB Administration"](#): Describes the display that allows you to modify your Service Data Model.

CMDB Administration

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 right, it displays the date and time '07-May-2019 15:56' and a 'DATA' status 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, which is currently open. The main area of the dialog is titled 'Show items with value that:' and contains two 'Contains' dropdown menus, each followed by a text input field. Below these are 'And' and 'Or' dropdown menus. 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.

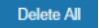
Criticality is the importance level of a CI to your organization. Criticality values range from **A** to **E**, where **A** is the highest Criticality and **E** is the lowest Criticality (with equally spaced intermediate values). This value is used to calculate **Alert Impact** (the maximum Alert Severity 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

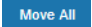
When you move a Service, Group or Area (Owners cannot be moved) you move it up one level in the CMDB and all lower level CMDB associations (Services, Groups and Areas and associated CIs) go with it.

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:

Owner	Select an Owner to filter by. The Owner selected populates the Area , Group and Service drop-down menus.
Manage Owner	<p>Opens a dialog that enables you to Delete, Rename or Merge the Owner.</p> <p>Delete removes the Owner from the CMDB database as well as all CMDB data and CIs associated with the Owner.</p> <p>Rename Changes all records for the Owner to a new name. Rename is disabled when the name you are typing in the text box already exists in the CMDB.</p> <p>Merge Changes all records for the Owner 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.</p> <p>Note: You cannot move Owners.</p>
Area	Select an Area to filter by. The Area selected populates the Group and Service drop-down menus.
Manage Area	<p>Opens a dialog that enables you to Delete, Rename or Merge the Area.</p> <p>Delete removes the Area from the CMDB database as well as all CMDB data and CIs associated with the Area.</p> <p>Rename Changes all records for the Area to a new name. Rename is disabled when the name you are typing in the text box already exists in the CMDB.</p> <p>Merge Changes all records for the Area 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.</p> <p>Move Changes all records for the Area to a different, already existing name in the CMDB that you choose from the New Area drop-down menu.</p>
Group	Select a Group to filter by. The Group selected populates the Service 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 CIs...** to open the **Find CIs 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.

Custom Tab

The **CUSTOM** tab provides a location where you can add your own custom tab, views and displays.

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.

Heatmaps

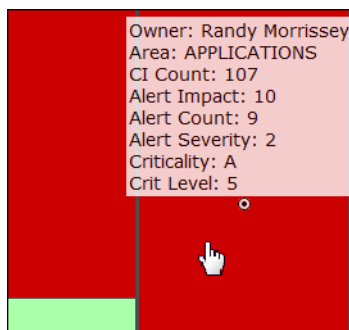
Heatmaps organize CIs (according to the Service Data Model) into rectangles and use color to highlight the most critical value in each. Heatmaps enable you to view various alert metrics in the same heatmap using drop-down menus. Each Metric has a color gradient bar that maps relative values to colors.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

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**.



Tables

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

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"](#)

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.

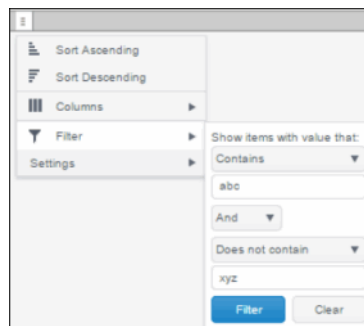
Column visibility changes are NOT reflected in an export to HTML and Excel.

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, (=, !=, >, >=, <, <=). 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.




217	emreference	sl.rtvew.sql.sqldb	RTVIEWFORJ1rootmy-secret-pw-joe.mysql/102
229	emreference	sl.rtvew.properties.queryTimeOut	\$rtvConfigDataServer:CONFIG_SERVER
216	emreference	sl.rtvew.sql.sqldb	10
			ALERTDEFS --- _none ---

Page 1 of 2 1 - 200 of 235 items

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.



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.

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.

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 MQ"](#)
- ["IBM WebSphere"](#)

Note: This document assumes familiarity with the products monitored. For additional details, refer to vendor documentation.

IBM DB2

The IBM DB2 displays provide extensive visibility into the health and performance of IBM DB2 databases. This features an ["IBM DB2 Overview Display"](#) (pictured below), and the following Views which can be found under **Components** tab > **Databases**> **IBM DB2 Database**:

- ["DB2 Instances View"](#): These displays present metrics about IBM DB2 instances.
- ["DB2 Members View"](#): These displays present metrics about IBM DB2 members.
- ["DB2 Databases View"](#): These displays present metrics about IBM DB2 databases.

IBM DB2 Overview Display

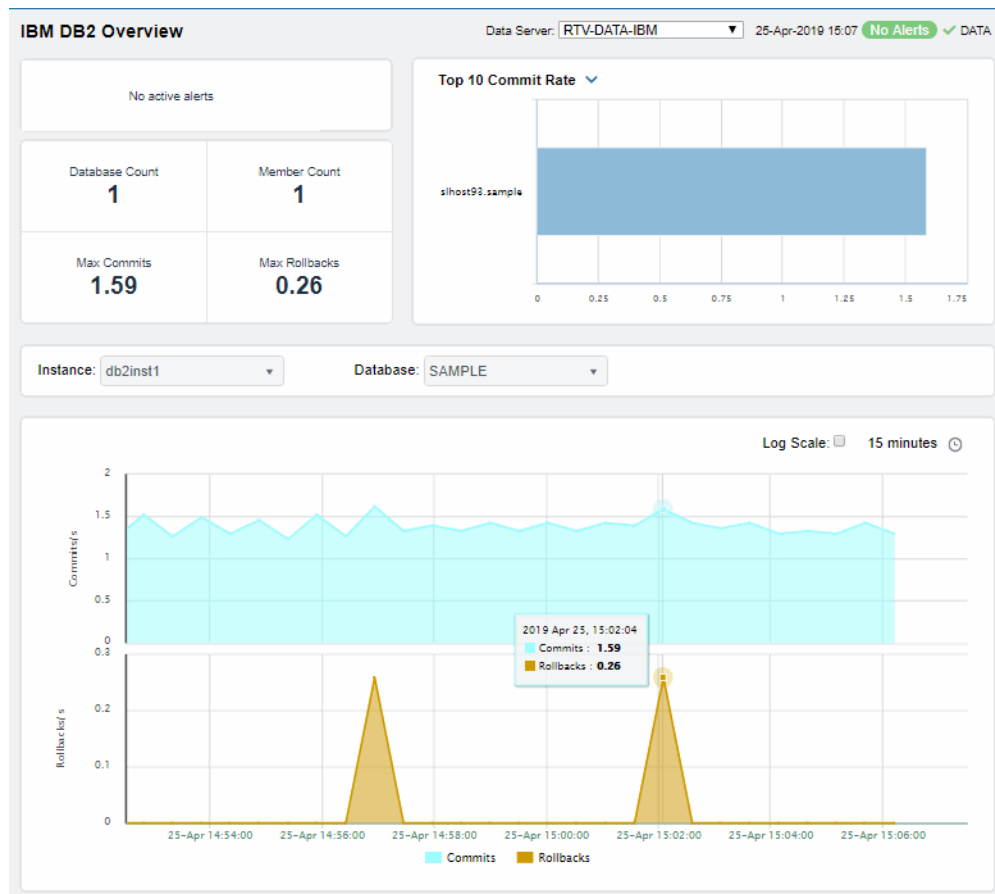
The **IBM DB2 Database Overview** is the top-level display for the IBM DB2 Database Solution Package, which provides a good starting point for immediately getting the status of all your IBM DB2 databases on your Data Server.

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 graph 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”](#), 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

Displays in this View are:

- [“DB2 Instances Table”](#)
- [“DB2 Instances Heatmap”](#)
- [“Instance Summary”](#)

DB2 Instances Table

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”](#) 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

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:

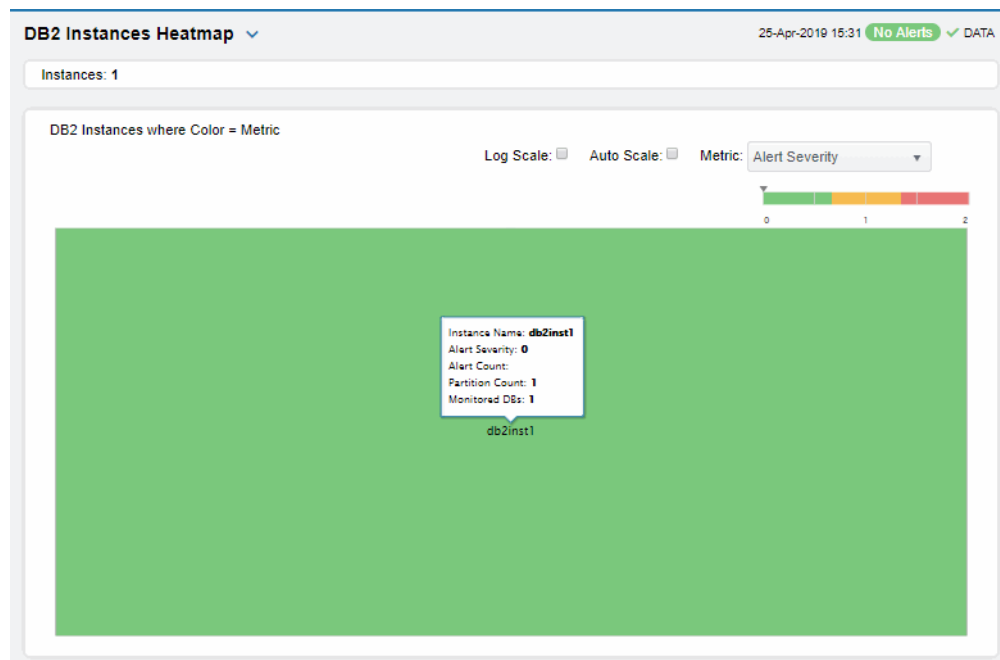
- (2) Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
- (1) Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
- (0) Green indicates that no metrics have exceeded their alert thresholds.

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](#)” display.



Instance Summary

Select an **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”](#) 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”](#) display.

The screenshot displays the 'DB2 Instance Summary' dashboard. At the top, it shows the instance name 'db2inst1' and 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' table lists one host, SLHOST93, with details on memory, CPUs, and CPU loads. The 'Databases' table shows one database, SAMPLE, with various performance metrics. 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.

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.

Instance Databases Table

Each row is a different database. 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.
Avg Lock Escalations per Activity	The average number of application deadlock escalations per activity.
Avg Lock Timeouts per Activity	The average number of application deadlock timeouts per activity.
Avg Lock Waits per Activity	The average number of application deadlock waits per activity.
Rows Read per Rows Returned	The number of rows read per number of rows returned.
Activities/sec	The number of activities per second.
App Requests/sec	The number of application requests per second.

Commits/sec	The number of application commits per second.
Rollbacks/sec	The number of application rollbacks per second.
Buffer Pool 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.
Activity Wait Time %	The percentage wait time taken by activities.
Avg Request CPU Time	The average amount of CPU time used by requests, in seconds.
Compile Proc Time %	The percentage of time used for compiling processes.
Routine Time Request %	The percentage of time used for routine request processes.
Section Time %	The percentage of time used for section processes.
Section Sort Time %	The percentage of time used for sorting section processes.
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.
Transaction Time %	The percentage of time used for transaction processes.
Utils Proc Time %	The percentage of time used for utilities processes.
Timestamp	The data and time this data was last updated.

DB2 Members View

Displays in this View are:

- [“DB2 Members Table”](#): Get configuration details and utilization metrics for host members of an instance.
- [“DB2 Member Summary”](#): View trend graph for CPU load and CPU utilization for each member of an instance.

DB2 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 “[DB2 Member Summary](#)” display.

DB2 Members Table 25-Apr-2019 15:36 DATA

Instance: db2inst1 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).

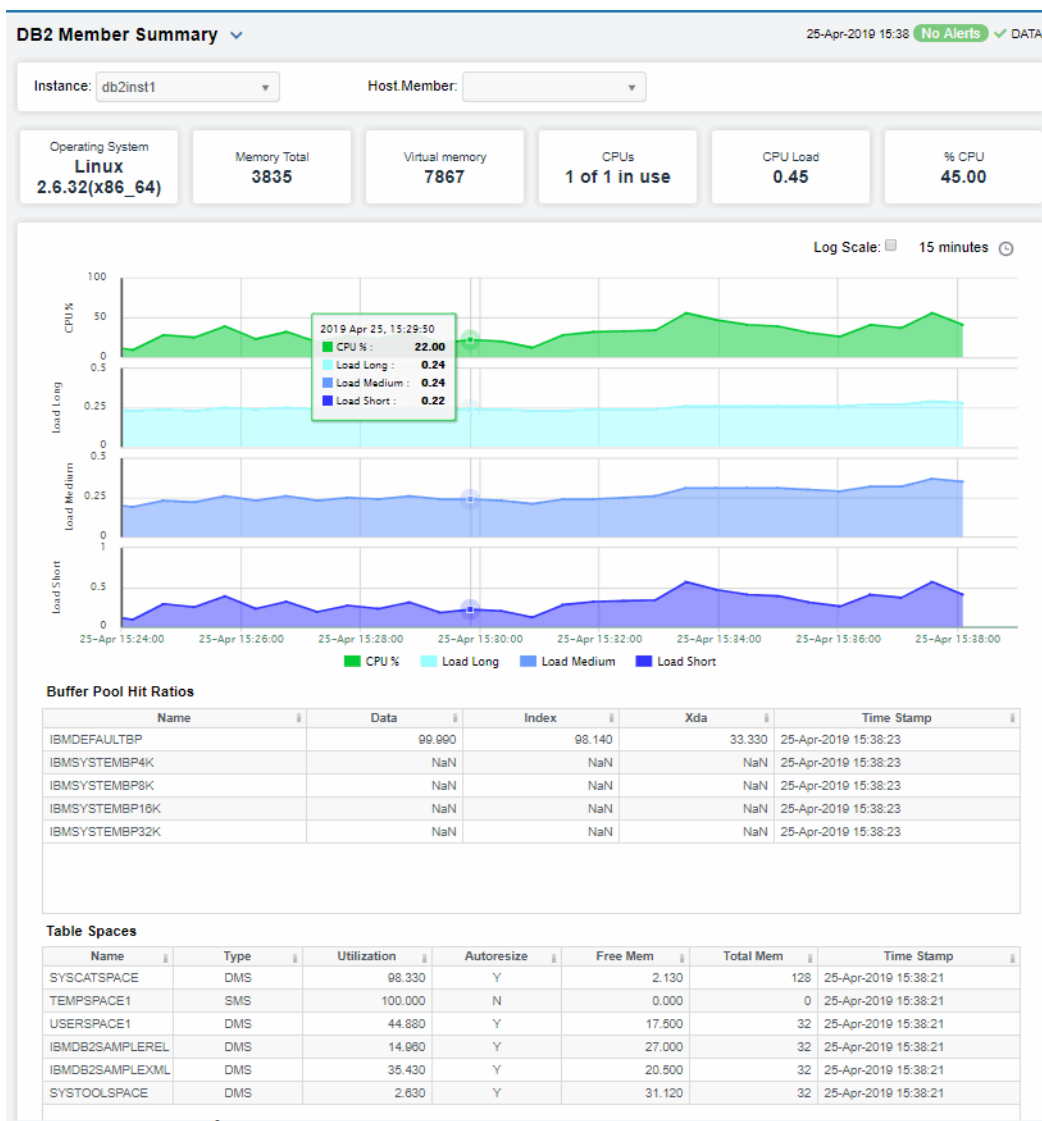
- CPU Load Long** Percentage of CPU load over the long term (defined by the IBM DB2 system, for example, 10-15 minutes).
- % CPU Usage** The percentage of CPU used.
- Timestamp** The data and time this data was last updated.

DB2 Member Summary

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.

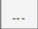
Auto Extend Indicates whether auto extend is enabled. Y/N


Trend Graph


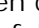
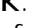
For the selected host member, trend graphs trace as follows:

- **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

Displays in this View are:

- ["DB2 Databases Table"](#)
- ["Single DB2 Database Summary"](#)
- ["Partition Table"](#)
- ["Partition Heatmap"](#)

DB2 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 “Single DB2 Database Summary” 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.

Avg Lock Escalations per Activity	The average number of application deadlock escalations per activity.
Avg Lock Timeouts per Activity	The average number of application deadlock timeouts per activity.
Avg Lock Waits per Activity	The average number of application deadlock waits per activity.
Rows Read per Rows Returned	The number of rows read per number of rows returned.
Activities/sec	The number of activities per second.
App Requests/sec	The number of application requests per second.
Commits/sec	The number of application commits per second.
Rollbacks/sec	The number of application rollbacks per second.
Buffer Pool 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.
Activity Wait Time %	The percentage wait time taken by activities.
Avg Request CPU Time	The average amount of CPU time used by requests, in seconds.
Compile Proc Time %	The percentage of time used for compiling processes.
Routine Time Request %	The percentage of time used for routine request processes.
Section Time %	The percentage of time used for section processes.
Section Sort Time %	The percentage of time used for sorting section processes.
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.
Transaction Time %	The percentage of time used for transaction processes.
Utils Proc Time %	The percentage of time used for utilities processes.
Activity Time (ms)	The amount of time, in milliseconds, used for activity processes.
Agent Idle Time (ms)	The amount of time, in milliseconds, that agents were in an idle state.
CPU Time (ms)	The amount of time, in milliseconds, used for CPU processes.
Processing Time (ms)	The amount of time, in milliseconds, used for processing requests.
Request Time (ms)	The amount of time, in milliseconds, used for processing requests.
DB Processing %	The percentage of time used for database processes.
DB Wait %	The percentage of time used for database waits.
# Connections	The current number of database connections.
Timestamp	The date 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

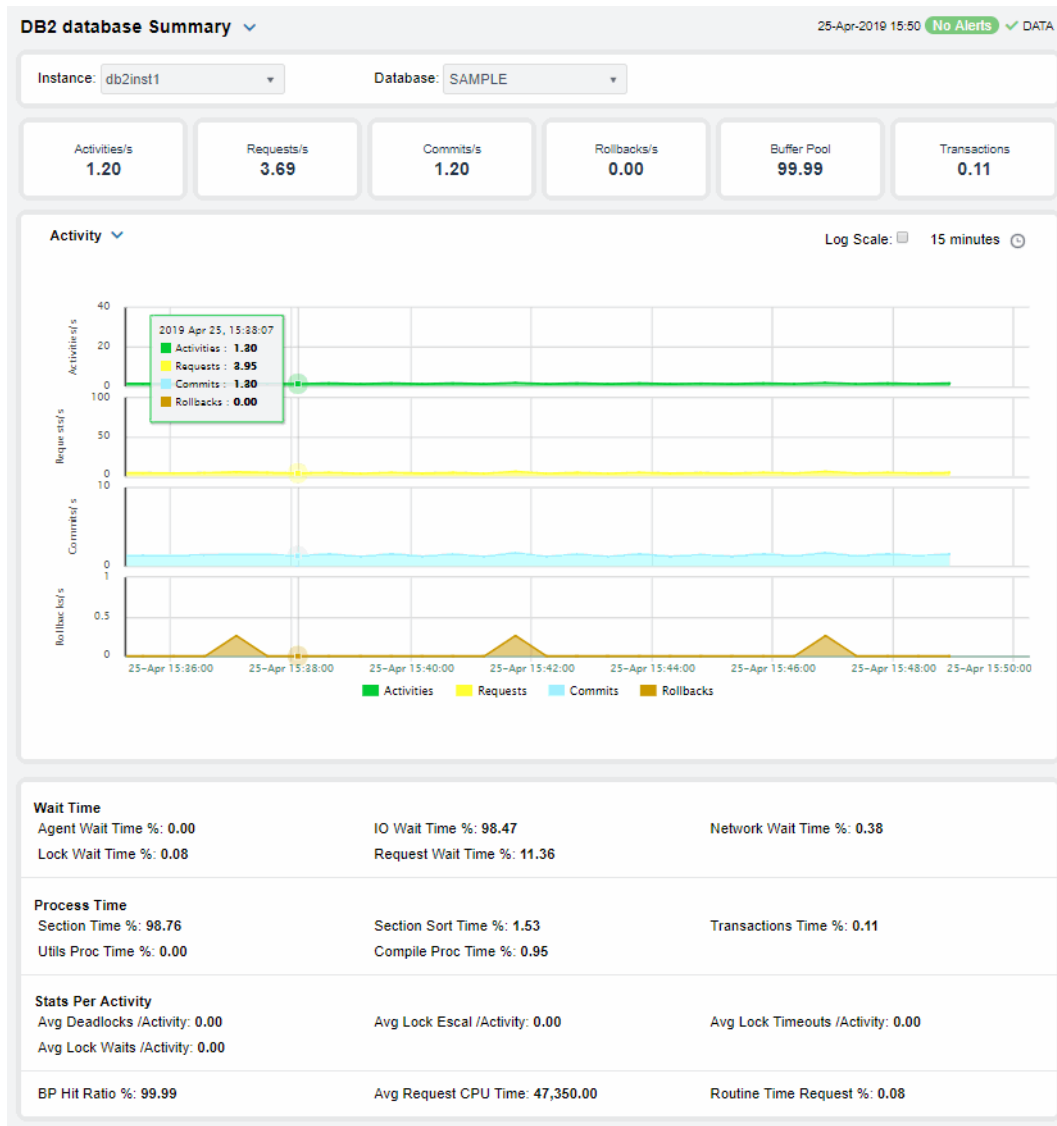
Single DB2 Database Summary

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

Activities	Total number of activities in the last 1 second.	Value per second.
App Requests	Total number of application requests in the last 1 second.	Value per second.
Commits	Total number of commits in the last 1 second.	Value per second.
Rollbacks	Total number of rollbacks in the last 1 second.	Value per second.

Avg # per Activity

Deadlocks	The average number of deadlocks per activity in the last 1 second.
Lock Escal	The average number of lock escalations per activity in the last 1 second.
Lock Timeouts	The average number of lock timeouts per activity in the last 1 second.
Lock Waits	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.


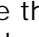
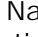
Overall Execution Trends

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.

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.

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 row will be deleted if no data is received within the time specified for deletion.

DB2 Status

The current DB2 status. For example, ACTIVE.

DB Status

The current database status. For example, ACTIVE.

DB Start Time

The date and time the database was last started.

Local Cons

The number of local connections.

Remote Cons

The number of remote connections.

App Cur Cons

The number of currently connected applications.

SQL Faults/sec

The number of SQL faults per second.

Rows Read /sec

The number of rows read per second.

Rows Selected /sec	The number of rows selected per second.
Rows Changed /sec	The number of rows changed per second.
SQL Selects /sec	The number of SQL selects per second.
Commits /sec	The number of commits per second.
Rollbacks /sec	The number of rollbacks per second.
Update/Del/Ins/Stmts / sec	The number of updates, deletions, insertions and statements per second.
Avg Sort Time/Transaction	The average amount of time for sorting transactions.
Product Name	The name of the product.
Service Level	The service level for the product.
Pool Data 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.
Pool TmpData Hit Ratio	Refer to vendor documentation for details.
Pool TmpIndex Hit Ratio	Refer to vendor documentation for details.
Pkg Cache Inserts/K-Trans	Refer to vendor documentation for details.
Lock Wait Time/K-Trans	Refer to vendor documentation for details.
Dirty Steal Triggers/K-Trans	Refer to vendor documentation for details.
Deadlocks & Lock Timeouts/K-Tran	Refer to vendor documentation for details.
Avg Log Write Time / Trans	Refer to vendor documentation for details.
% Agent Usage	The percentage used by agents.
Agents Registered (Top)	The number of registered agents.
Rows Read Returned Ratio	The number of rows read per number of rows returned.
Select %	Refer to vendor documentation for details.
Phys Buffer Pool Read Ratio	Refer to vendor documentation for details.
Phys Buffer Pool Write Ratio	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?

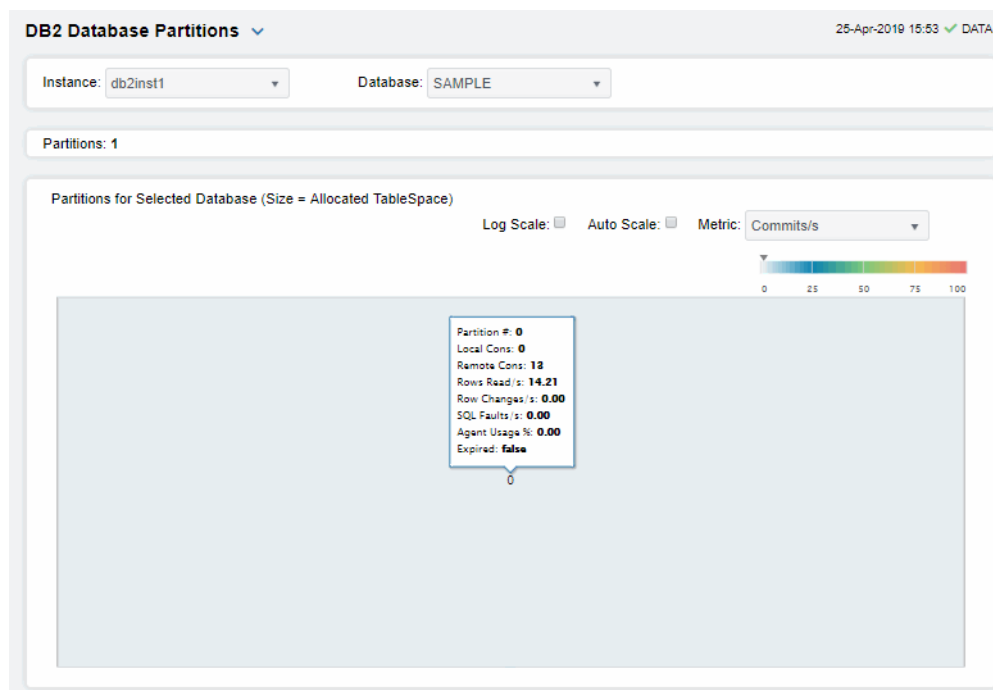
Each rectangle in the heatmap represents a different partition, where the rectangle color indicates the most critical alert state for items associated with that partition, and the rectangle size represents the tablespace allocation size for the partition.

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 “[Single DB2 Database Summary](#)” 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.

Rows Changed /sec	<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 0 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">  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 Select Stmts/sec	<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 0 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">  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.
Rollbacks/sec	<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">  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 Update/Del/Ins/Stmts /sec	<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">  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.

IBM MQ

The IBM MQ HTML displays provide extensive visibility into the health and performance of IBM MQ. This features an ["IBM MQ Overview Display"](#) (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"**: The displays in this View present performance and utilization metrics for your IBM MQ Brokers.
- **"IBM MQ Channels View"**: The displays in this View present performance and utilization metrics for your IBM MQ Channels.
- **"IBM MQ Queues View"**: The displays in this View present performance and utilization metrics for your IBM MQ Queue Managers.
- **"IBM MQ Topics View"**: The displays in this View present performance and utilization metrics for your IBM MQ Topics.
- **"IBM MQ Topic Configs View"**: The displays in this View present performance and utilization metrics for your IBM MQ Topic Configs.
- **"IBM MQ Topic Publishers View"**: The displays in this View present performance and utilization metrics for your IBM MQ Topic Publishers.
- **"IBM MQ Topic Subscribers View"**: The displays in this View present performance and utilization metrics for your IBM MQ Topic Subscribers.
- **"IBM MQ Topic Publisher Totals View"**: The displays in this View present performance and utilization metrics for your IBM MQ Topic Publisher Totals.
- **"IBM MQ Topic Subscriber Totals View"**: The displays in this View present performance and utilization metrics for your IBM MQ Topic Subscriber Totals.
- **"IBM MQ Subscriptions View"**: The displays in this View present performance and utilization metrics for your IBM MQ Subscriptions.
- **"IBM MQ Subscription Configs View"**: The displays in this View present performance and utilization metrics for your IBM MQ Subscription Configs.

IBM MQ Overview Display

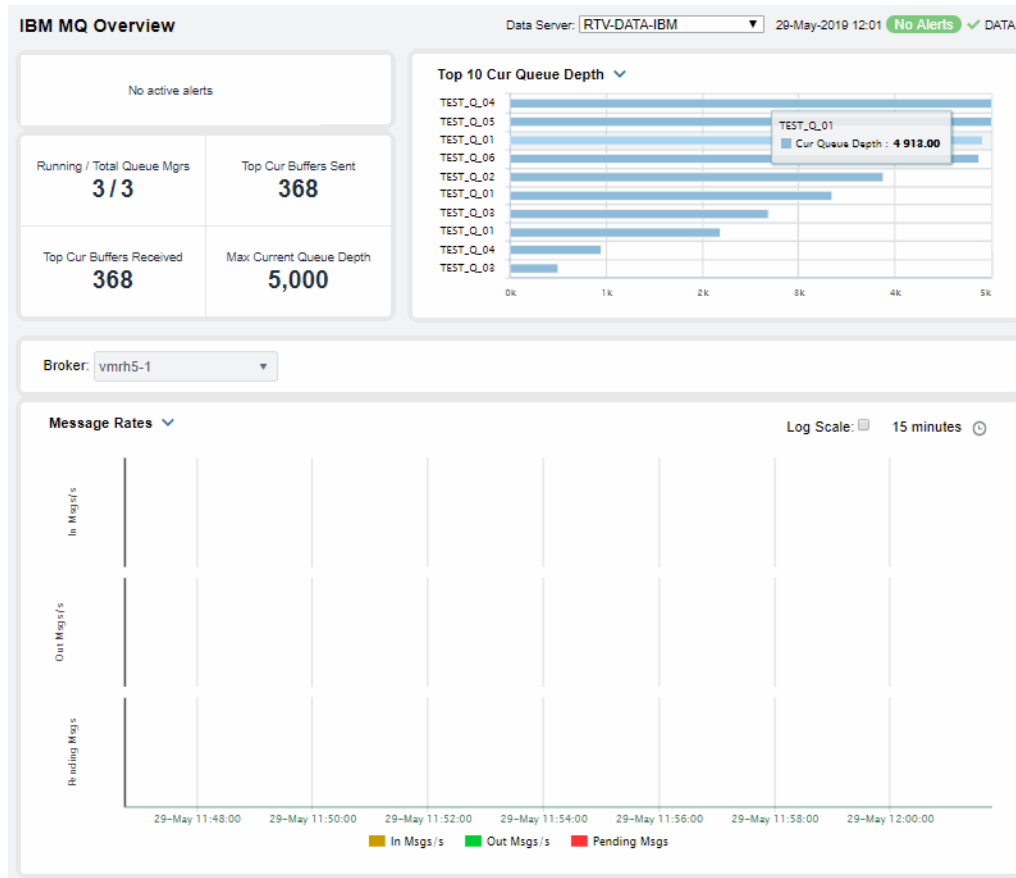
The **IBM MQ Overview** is the top-level display for the IBM MQ Solution Package, which provides a good starting point for immediately getting the status of all IBM MQ objects on your Data Server.

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 Recieved**.
- 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"**.

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"** 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

See performance and utilization metrics for all of your IBM MQ Brokers.

Displays in this View are:

- [“IBM MQ Broker Connections Table”](#): View a sortable list of utilization metrics for all IBM MQ brokers and compare broker metrics.
- [“IBM MQ Brokers Heatmap”](#): This display presents heatmap view of IBM MQ brokers and their alert states.
- [“IBM MQ Broker Summary”](#): This display presents performance metrics for a single IBM MQ Broker, as well as detailed metrics for its channels and queues.

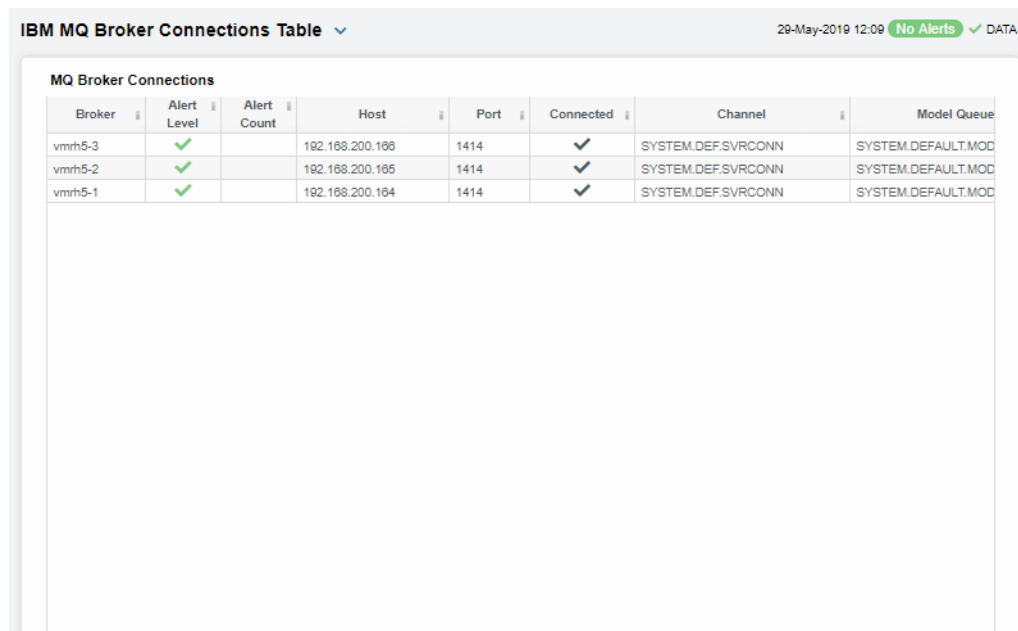
IBM MQ Broker Connections Table

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](#)” display.



Broker	Alert Level	Alert Count	Host	Port	Connected	Channel	Model Queue
vmrh5-3	✓		192.168.200.166	1414	✓	SYSTEM.DEF.SVRCONN	SYSTEM.DEFAULT.MOD
vmrh5-2	✓		192.168.200.165	1414	✓	SYSTEM.DEF.SVRCONN	SYSTEM.DEFAULT.MOD
vmrh5-1	✓		192.168.200.164	1414	✓	SYSTEM.DEF.SVRCONN	SYSTEM.DEFAULT.MOD

MQ Broker Connections Table

Each table row is a different connection. Column values describe the connection except where noted.


Broker	The name of the broker.
Alert Level	The current alert severity: <ul style="list-style-type: none"> ● 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.
Host	The name of the host.
Port	The port number used.




Connected	When checked, denotes that the broker is connected.
Channel	The name of the channel.
Model Queue Name	Named model queue of the connection.
Max Retries	Maximum number of subsequent connection retry attempts.
Retry Interval	Minimum interval (in seconds) between connection retry attempts.
Wait Interval	Wait interval (in seconds) between attempts to create a connection.
time_stamp	The data and time of the last data update.

IBM MQ Brokers Heatmap

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:

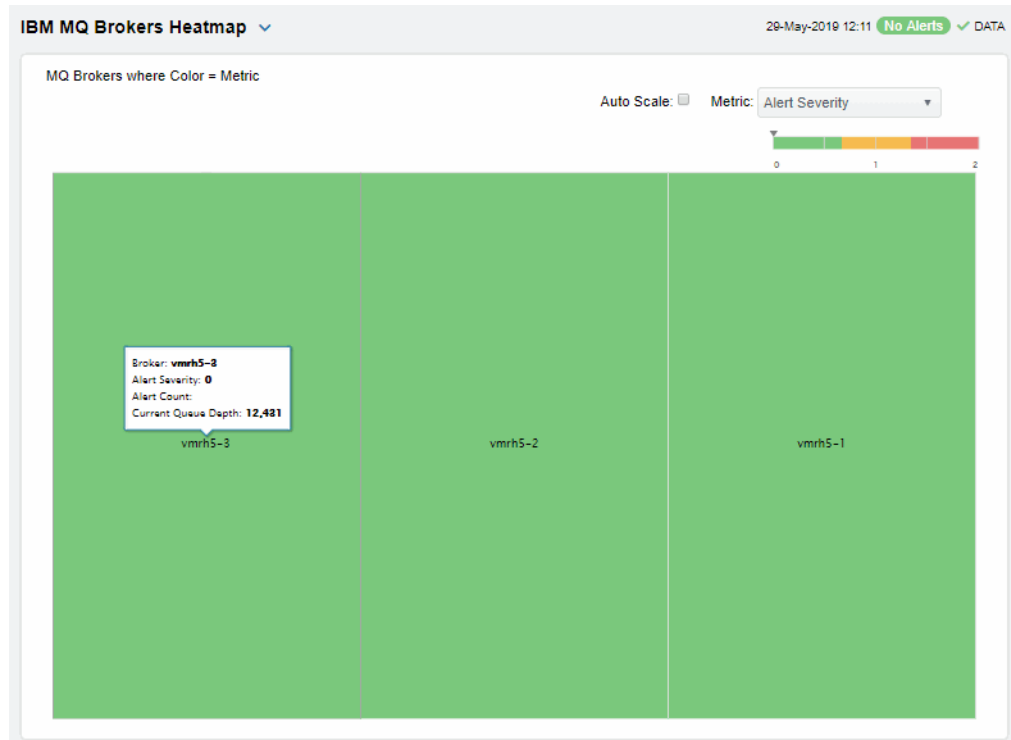
-  **(2)** Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
-  **(1)** Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
-  **(0)** Green indicates that no metrics have exceeded their alert thresholds.

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"](#) 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.



Fields and 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 Scale** is not selected.



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.
Current Queue Depth	The current queue depth. The color 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 Level for the MqBrokerQueueDepthHigh alert in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

IBM MQ Broker Summary

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"](#) 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”** 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.




Filter By:

MQ Broker Select the broker for which you want to view data.

Fields and Data




Running Queue Mgrs	The number of queue managers running on the broker.
Channels	The current number of channels on the broker.
Queues	The current number of queues on the broker.
Total Queue Depth	The total queue depth across all queues on the broker.
Total Open Inputs	The total number of open inputs across all queues on the broker.
Total Open Outputs	The total number of open outputs across all queues on the broker.
Total Current Queue Depth Trend Graph	Current Queue Depth -- Traces the total current queue depth for the selected 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	Select a time range from the drop down menu varying from 5 Minutes to Last 7 Days . By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

MQ Channels for Selected Broker Table

Channel Name	The name of the channel.
---------------------	--------------------------

Type	The type of channel.
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.
Rcvd Buffers	The number of buffers received.
Current Sent Buffers	The number of buffers received since the last data update.
Current Rcvd Bytes	The number of bytes received since the last data update.
Current Completed	The number of batches completed since the last data update.
MQCACH_MCA_USER_ID	The user ID used by the MCA.
MQCACH_RCV_EXIT_NAME	The receive exit name.
MQCACH_RCV_EXIT_USER_DATA	The receive exit user data.
MQCACH_SECURITY_EXIT_NAME	The security exit name.
Expired	When checked, performance data has not been received in the time specified in the Duration region on the RTView Configuration Application > Solution Package Configuration > IBM MQ > DATA STORAGE tab.
Critical/Warning	The total number of critical and warning alerts.
Connected	When true , denotes that the broker is connected.
Last Update	The date and time of the last data update.

IBM MQ Channels View

See performance and utilization metrics for all of your IBM MQ Brokers.

Displays in this View are:

- **"IBM MQ Channels Table"**: This display presents a high-level perspective of utilization metrics for each IBM MQ Broker.
- **"IBM MQ Channels Heatmap"**: This display presents detailed performance metrics for each channel
- **"IBM MQ Channel Summary"**: This display presents additional configuration metrics for a single channel.

IBM MQ Channels Table

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"** display.

Note: This display contains vendor data. Refer to vendor documentation for details.

IBM MQ Channels Table 29-May-2019 12:20 No Alerts DATA

Broker: - All -




Channels: 4

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

MQ Broker	Select the broker for which you want to view data, or select All to view data for all brokers.
Channels	The total number of channels found on the broker and listed in the table.

Table

Each table row is a different channel on the selected broker. Column values describe the channel.

Broker	The name of the broker.
Channel Name	The name of the channel.
Type	The type of channel.
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 active alerts on the channel.
Rcvd Buffers	The number of buffers received.
Current Rcvd Buffers	The number of buffers received since the last data update.
Sent Buffers	The number of buffers sent.
Current Sent Buffers	The number of buffers sent since the last data update.
Rcvd Bytes	The number of bytes received.
Current Rcvd Bytes	The number of bytes received since the last data update.
Sent Bytes	The number of bytes sent.
Current Sent Bytes	The number of bytes sent since the last data update.
Completed Batches	The number of batches completed.
Current Completed Batches	The number of batches completed since the last data update.
Description	A textual description of the channel.
MQCACHE_MCA_USER-ID	The user ID used by the MCA.
MQCACHE_RCV_EXIT_NAME	Denotes the name of the user exit program that was run by the channel receive user exit.

MQCACHE_RCV_EXIT_USER_DATA	The user data that is passed to the receive exit.
MQCACHE_SEX_EXIT_NAME	Denotes the name of the exit program that was run by the channel security exit.
MQCACHE_SEX_EXIT_USER_DATA	The user data that is passed to the security exit.
MQCACHE_SEND_EXIT_NAME	The send exit name.
MQCACHE_SEND_EXIT_USER_DATA	The user data that is passed to the send exit.
MQCACHE_SSL_CIPHER_SPEC	Denotes the single CipherSpec for a TLS or SSL connection.
MQCACHE_SSL_PEER_NAME	The Distinguished Name (DN) of the certificate from the peer queue manager or client at the other end of a IBM WebSphere MQ channel.
MQCA Alteration Date	The date the MQ CA was last modified.
MQCA Alteration Time	The time the MQ CA was last modified.
IACH HB Interval	The ACH heartbeat interval setting.
IACH Hdr Compression	The ACH header data compression techniques supported by the channel.
IACH Keep Alive Internal	The ACH keep alive interval setting (the timeout value for the channel).
IACH Max Msg Length	The ACH maximum message length setting.
MQIACH_MSG_COMPRESSION	The ACH message data compression techniques supported by the channel.
MQIACH_SSL_CLIENT_AUTH	Denotes whether the channel needs to receive and authenticate an SSL certificate from an SSL client.
MQIACH_XMIT_PROTOCOL_TYPE	The transport (transmission protocol) type used.
MQIA_MONITORING_CHANNEL	Denotes the attribute used to control the collection of online monitoring data.


Max Message Length	The maximum length of messages on the channel.
Status	The channel status.
Transmission Queue Name	The name of the queue that transmits for the channel.
Connection Name	The name of the connection.
Expired	When checked, performance data has not been received in the time specified in the Duration region on the RTView Configuration Application > Solution Package Configuration > IBM MQ > DATA STORAGE tab.
Time Stamp	The date and time of the last data update.

IBM MQ Channels Heatmap

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:

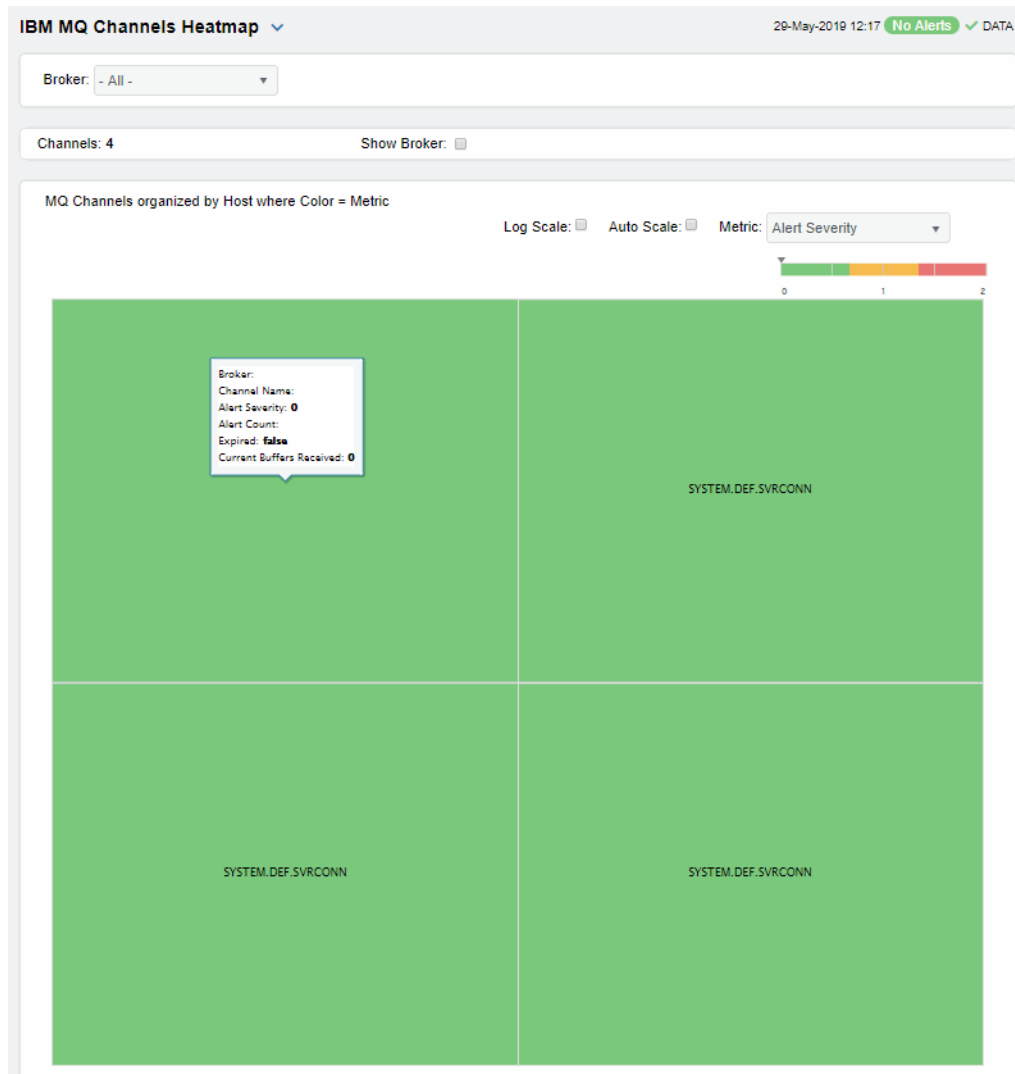
- **(2)** Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
- **(1)** Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
- **(0)** Green indicates that no metrics have exceeded their alert thresholds.

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"](#) 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.



Filter By






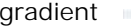



Broker Select the broker containing the channels for which you want to view data, or select **All** to view all channels for all brokers.

Fields and Data

Channels The number of channels found on the broker(s) and listed in the heatmap.

Show Broker Select this check box to display the name of the broker in the heatmap.

Heatmap

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.
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 Scale is not selected.
Metric	Choose a metric to view in the display. For details about the data, refer to vendor documentation.
Alert Severity	<p>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:</p> <ul style="list-style-type: none">  (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.
Current Buffers Received	The current number of buffers 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 maximum number of buffers received in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
Current Buffers Sent	The current number of buffers 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 maximum number of buffers sent in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
Current Bytes Received	The current 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 0 to the defined Alert Level for the MqChannelBytesInHigh alert in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
Current Bytes Sent	The current 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 0 to the defined Alert Level for the MqChannelBytesOutHigh alert in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

IBM MQ Channel Summary

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.

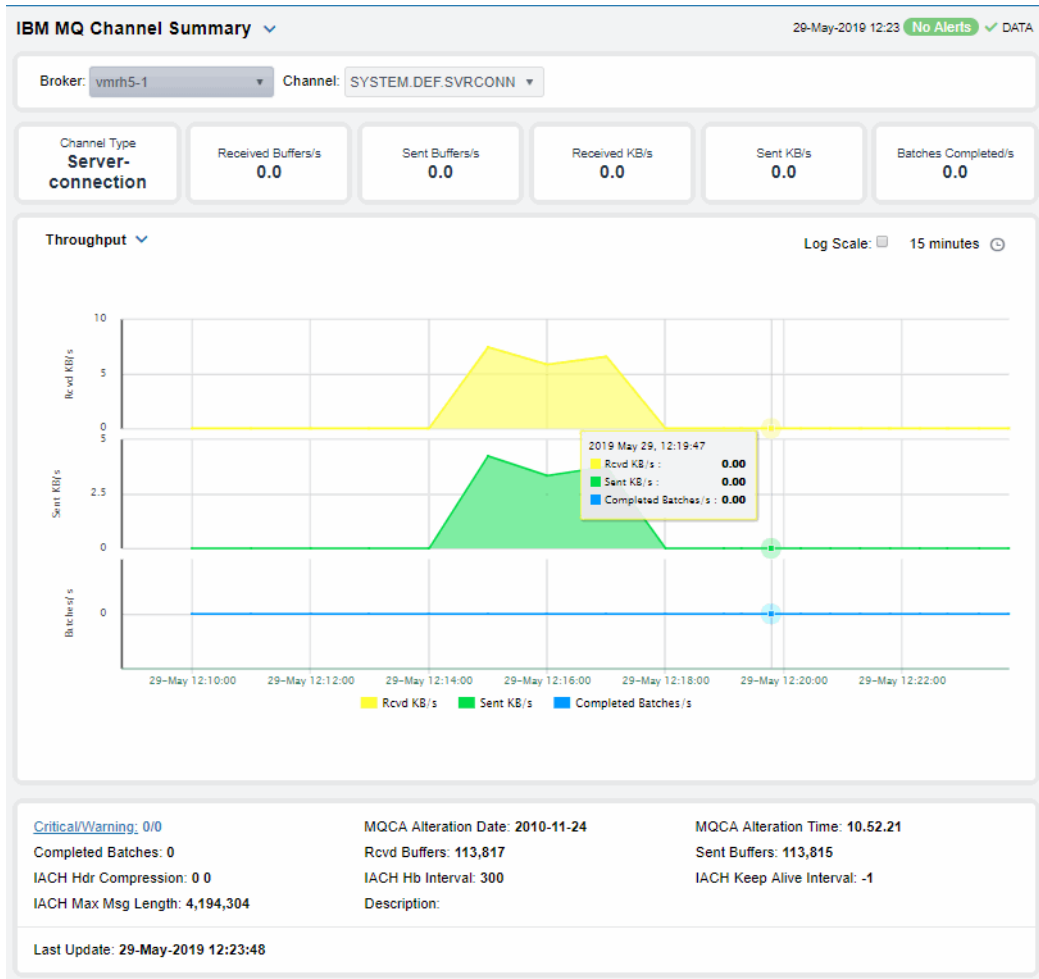
Choose a broker and a channel from the drop-down menus. Clicking on the information boxes at the top of the display (such as **Received Buffers/second**, **Sent Buffers/second** and **Batches Completed**) takes you to the ["IBM MQ Channels Table"](#) display, where you can sort and compare the performance 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.



Filter By

- MQ Broker** Select the broker containing the channel for which you want to view data.
- Channel** Select the channel for which you want to view data.

Fields and Data

- Channel Type** The type of channel.
- Received Buffers/s** The rate of buffers received, per second, by this channel.
- Sent Buffers/s** The rate of buffers sent, per second, by this channel.
- Received KB/s** The rate of kilobytes received, per second, by this channel.
- Sent KB/s** The rate of kilobytes sent, per second, by this channel.

	Batches Completed/s	The rate of batches being completed, per second, by this channel.
Trend Graphs	Throughput	
	Rcvd KBytes/sec:	The current number of kilobytes received per second.
	Sent KBytes/sec:	The current number of kilobytes sent per second.
	Completed Batches/sec:	The current number batches completed per second.
	Buffer Flow	
	Rcvd Buffers/sec:	The current number of buffers received per second.
	Sent Buffers/sec:	The current number of buffers sent per second.
	Completed Batches/sec:	The current number batches completed per second.
Critical/Warning		The total number of critical and warning alerts.
Completed Batches		The total number of completed batches.
IACH Hdr Compression		The ACH header data compression techniques supported by the channel.
IACH Max Msg Length		The ACH maximum message length setting.
MQCA Alteration Date		The date the MQ CA was last modified.
Rcvd Buffers		The total number of received buffers.
IACH Hb Interval		The ACH heartbeat interval setting.
Description		A textual description of the channel.
MQCA Alteration Time		The time the MQ CA was last modified.
Sent Buffers		The total number of sent buffers.
IACH Keep Alive Interval		The ACH keep alive interval setting (the timeout value for the channel).
Last Update		The date and time of the last data update.

IBM MQ Queues View

See performance and utilization metrics for all of your IBM MQ queue managers.

Displays in this View are:

- [“IBM MQ Queue Managers Table”](#): This display presents a high-level perspective of utilization metrics for each IBM MQ queue managers.
- [“IBM MQ Queues Table”](#): This display lists all IBM MQ queues with detailed performance metrics and configuration information.
- [“IBM MQ Queues Heatmap”](#): This display presents a heatmap view of performance metrics and alert levels for one or all brokers.
- [“IBM MQ Queue Summary”](#): This display presents detailed performance metrics and configuration information for a single IBM MQ queue.

IBM MQ Queue Managers Table

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”](#) 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		

Filter By

Broker Select the broker for which you want to view data or select **All** to view data for all brokers.

Table Each table row is a different queue manager on the selected broker. Column values describe the queue.

Broker	The name of the broker.
Name	The name of the queue manager.
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 active alerts on the queue manager.
Status	The queue manager status (for example, Running).
Command Level	The command level.
Connection Count	The number of connections on the queue manager.

Dead letter queue	The number of undelivered messages in the dead letter queue.
Description	A textual description of the queue manager.
Max Message Length	The maximum message length sent or received by the queue manager.
Max priority	The queue manager rank in priority.
Platform	The queue manager platform type.
Host	The host name.
Connection	The connection name.
Expired	When checked, performance data has not been received in the time specified in the Duration region on the RTView Configuration Application > Solution Package Configuration > IBM MQ > DATA STORAGE tab.
Time Stamp	The date and time of the last data update.

IBM MQ Queues Table

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"](#) 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 PERS
TEST_Q_02	QM_0001	LOCAL	ONLINE	✓		0	0	3,880	5,000	NOT PERS
TEST_Q_03	QM_0001	LOCAL	ONLINE	✓		0	0	0	5,000	NOT PERS
TEST_Q_04	QM_0001	LOCAL	ONLINE	✓		0	0	947	5,000	NOT PERS
TEST_Q_05	QM_0001	LOCAL	ONLINE	✓		2	1	4,873	5,000	NOT PERS
TEST_Q_06	QM_0001	LOCAL	ONLINE	✓		0	0	4,887	5,000	NOT PERS
TEST_Q_07	QM_0001	LOCAL	ONLINE	✓		0	0	0	5,000	NOT PERS
TEST_Q_08	QM_0001	LOCAL	ONLINE	✓		0	0	0	5,000	NOT PERS
TEST_Q_09	QM_0001	LOCAL	ONLINE	✓		0	0	0	5,000	NOT PERS
TEST_Q_10	QM_0001	LOCAL	ONLINE	✓		0	0	0	5,000	NOT PERS
TEST_Q_01	QM_0001	LOCAL	ONLINE	✓		2	2	0	5,000	NOT PERS
TEST_Q_02	QM_0001	LOCAL	ONLINE	✓		0	0	0	5,000	NOT PERS
TEST_Q_03	QM_0001	LOCAL	ONLINE	✓		0	0	504	5,000	NOT PERS
TEST_Q_04	QM_0001	LOCAL	ONLINE	✓		0	0	0	5,000	NOT PERS
TEST_Q_05	QM_0001	LOCAL	ONLINE	✓		0	0	433	5,000	NOT PERS
TEST_Q_06	QM_0001	LOCAL	ONLINE	✓		0	0	0	5,000	NOT PERS
TEST_Q_07	QM_0001	LOCAL	ONLINE	✓		0	0	0	5,000	NOT PERS
TEST_Q_08	QM_0001	LOCAL	ONLINE	✓		0	0	0	5,000	NOT PERS
TEST_Q_09	QM_0001	LOCAL	ONLINE	✓		0	0	0	5,000	NOT PERS
TEST_Q_10	QM_0001	LOCAL	ONLINE	✓		0	0	0	5,000	NOT PERS
TEST_Q_01	QM_0001	LOCAL	ONLINE	✓		1	1	23	5,000	NOT PERS
TEST_Q_02	QM_0001	LOCAL	ONLINE	✓		0	1	0	5,000	NOT PERS
TEST_Q_03	QM_0001	LOCAL	ONLINE	✓		1	1	0	5,000	NOT PERS
TEST_Q_04	QM_0001	LOCAL	ONLINE	✓		1	0	5,000	5,000	NOT PERS
TEST_Q_05	QM_0001	LOCAL	ONLINE	✓		0	0	0	5,000	NOT PERS
TEST_Q_06	QM_0001	LOCAL	ONLINE	✓		2	3	0	5,000	NOT PERS
TEST_Q_07	QM_0001	LOCAL	ONLINE	✓		0	0	0	5,000	NOT PERS
TEST_Q_08	QM_0001	LOCAL	ONLINE	✓		0	0	0	5,000	NOT PERS
TEST_Q_09	QM_0001	LOCAL	ONLINE	✓		0	0	0	5,000	NOT PERS
TEST_Q_10	QM_0001	LOCAL	ONLINE	✓		0	0	0	5,000	NOT PERS

Filter By




MQ Broker Select the broker for which you want to view data or select **All** to view data for all brokers.

Queues The number of queuese found in the query and listed inthe table.

Table

Each table row is a different queue on the selected broker. Column values describe the queue.

- Broker** The name of the broker.
- 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).


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 on the queue.
Open Outputs	The number of outgoing transactions.
Open Inputs	The number of incoming transactions.
Cur Queue Depth	The current queue depth.
Max Q Depth	The maximum number of messages allowed on the queue at any one time.
Persistence	Denotes whether or not the queue manager is persistent (PERSISTENT/NOT PERSISTENT).
Description	The description of the purpose of the queue.
Max Msg Length	The maximum length of messages.
Host	The IP address of the host.
Default Priority	The default priority value for messages placed on the queue.
Get Messages	Denotes whether or not the queue enabled to get messages (GET ALLOWED/GET NOT ALLOWED).
Put Messages	Denotes whether or not the queue enabled to put messages (PUT ALLOWED/PUT NOT ALLOWED).
Scope	The defined scope setting for the queue.
Shareability	Denotes whether or not the queue is shareable (SHAREABLE/ NOT SHAREABLE).
Usage	The queue usage type (NORMAL/TRANSMISSION).
Connection	The name of the queue connection.
Expired	When checked, performance data has not been received in the time specified in the Duration region on the RTView Configuration Application > Solution Package Configuration > IBM MQ > DATA STORAGE tab.
Time Stamp	The date and time of the last data update.

IBM MQ Queues Heatmap

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 queue.

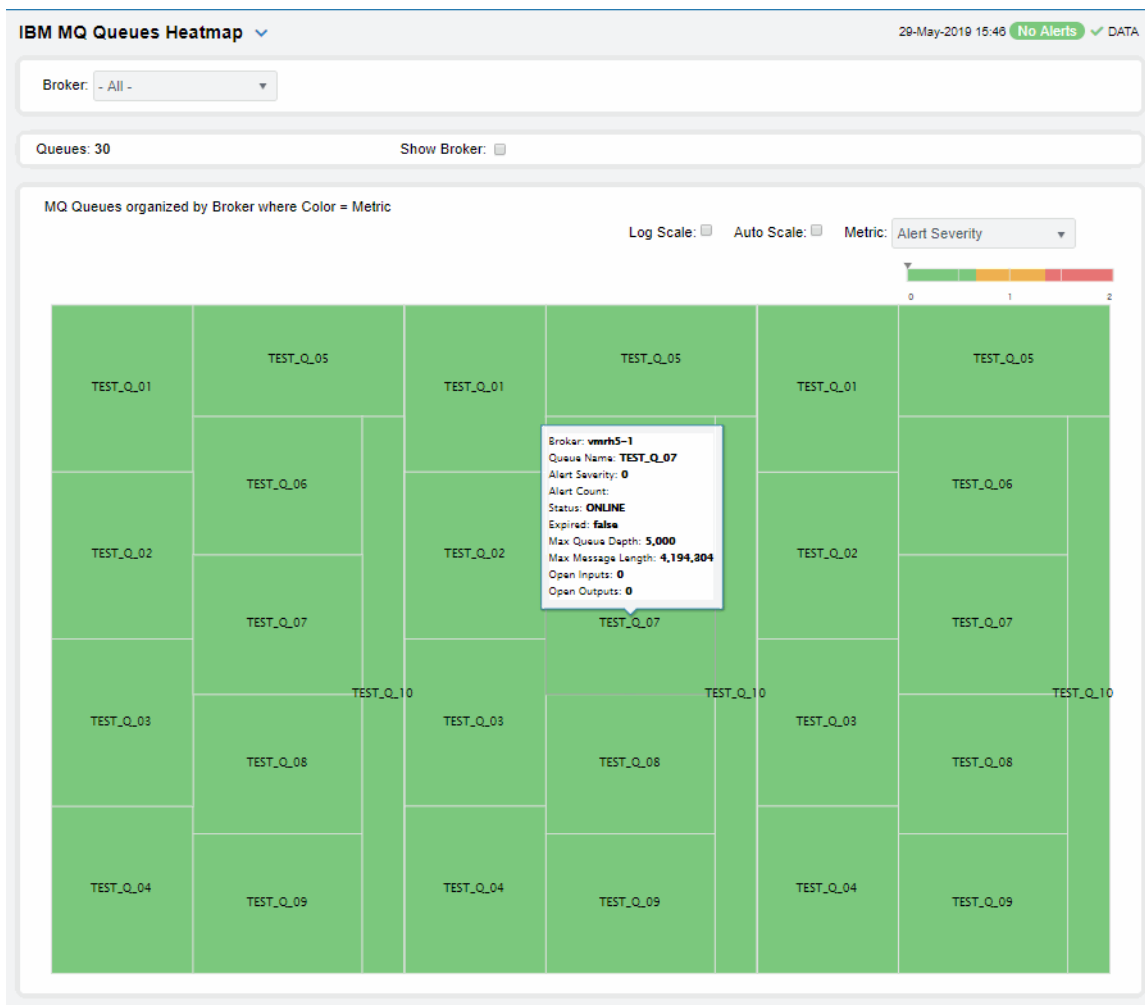
By default, the **Alert Severity** metric is shown. Values range from **0** - **2**, as indicated in the color gradient  bar:

- (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.

Investigate a queue by clicking a rectangle which opens the “[IBM MQ Queue Summary](#)” 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.



Filter By

Broker Select the broker containing the channels for which you want to view data, or select **All** to view all channels for all brokers.

Fields and Data

Queues The number of queues found on the broker(s) and listed in the heatmap.

Show Broker Select this check box to display the name of the broker in the heatmap.

Heatmap


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.


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 Scale** is not selected.


Metric Choose a metric to view in the display. For details about the data, refer to vendor documentation.

Alert Severity

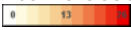
The current alert severity for items associated with the rectangle. Values range from 0 - 2, as indicated in the color gradient  bar, where 2 is the highest Alert Severity:

 (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.

Max Queue Depth

The maximum queue depth. The color 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 Level for the **MqQueueDepthHigh** alert in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

Max Message Length

The maximum message length. The color 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 message length in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

Open Inputs	The current number of open inputs. The color 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 open inputs in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
Open Outputs	The current number of open outputs. The color 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 open outputs in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

IBM MQ Queue Summary

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"](#) display, where you can sort and compare the performance values of all channels.

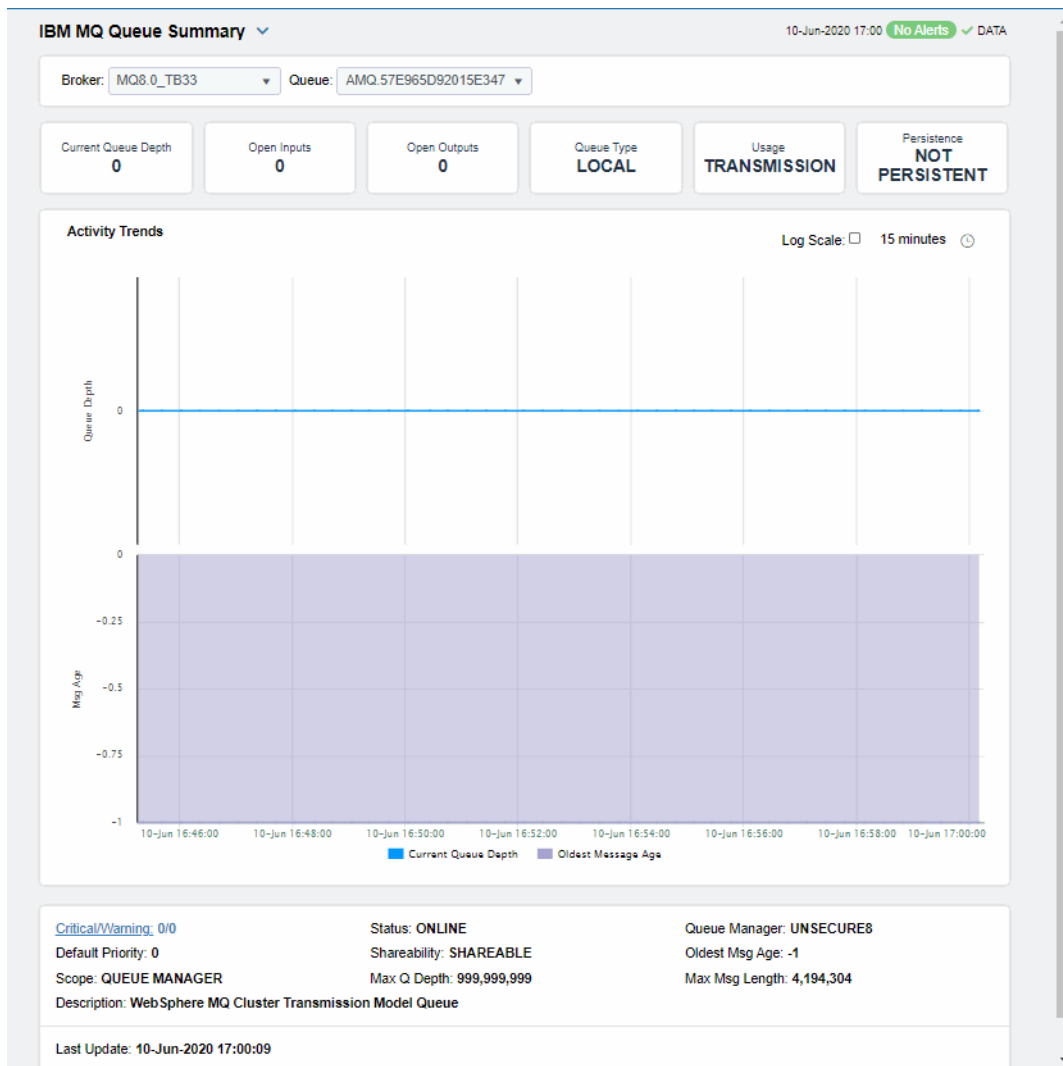
The trend graph traces the **Activity Trends (Current Queue Depth and Oldest Message Age)** rates for the selected queue. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your 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.

Queue Manager, **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.




Filter By

- MQ Broker** Select the broker containing the queue for which you would like to view data.
- Queue** Select the Queue for which you would like to view data.

Fields and Data

- Current Queue Depth** The current depth of the queue.
- Open Inputs** The total number of open inputs for the queue.
- Open Outputs** The total number of open outputs for the queue.

Queue Type	The type of queue.
Usage	The queue usage type (NORMAL/TRANSMISSION).
Persistence	Denotes whether or not the queue manager is persistent (PERSISTENT/NOT PERSISTENT).
Activity Trends Graph	<p>Current Queue Depth: Traces the current depth of the queue.</p> <p>Oldest Message Age: Traces the age of the oldest message in the queue.</p>
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	Select a time range from the drop down menu varying from 5 Minutes to Last 7 Days . By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

Critical/Warning	The total number of critical and warning alerts.
Shareability	Denotes whether or not the queue manager is shareable (SHAREABLE/NOT SHAREABLE).
Max Q Depth	The maximum number of messages allowed on the queue.
Description	The description of the queue.

Status	The status of the queue.
Oldest Msg Age	The age of the oldest message in the queue.
Max Msg Length	The maximum message length on the queue.
Default Priority	The default priority setting on the queue manager.
Scope	The defined scope for the queue.
Last Update	The data and time of the last data update.

IBM MQ Topics View

This View allows you to view performance and utilization metrics for all of your IBM MQ topics. Clicking **IBM MQ Topics** in the left menu opens the [“IBM MQ Topics Table”](#) display, which allows you to view current alert status and performance metrics of all topics on a specific MQ broker in a table format. Other displays available in this View are:

- [“IBM MQ Topics Heatmap”](#): Clicking **MQ Topics Heatmap** in the left/navigation menu opens this display, which allows you to view current alert status and performance metrics of all topics on a specific MQ broker in a heatmap format.
- [“IBM Topic Summary”](#): Clicking **MQ Topic Summary** in the left/navigation menu opens this display, which allows you to view current alert status and performance metrics for a specific topic on a specific MQ broker.

IBM MQ Topics Table

This display allows you to view current alert status and performance metrics of all topics on a specific MQ broker in a table format.

Each row in the table contains data for a particular topic. 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**.

You can also 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 topic by double-clicking a row, which opens the [“IBM Topic Summary”](#) display.

IBM MQ Topics Table 22-Oct-2019 10:49 No Alerts ✓ DATA

Broker: MQ7.5_TB23

Topics: 5

All Topics Table

Subscription Count	Admin Topic Name	Cluster Name	Cluster
0	SYSTEM.BROKER.ADMIN.STREAM		
0	QA.TEST.TOPIC		
0			
0			
1			

Filter By:

MQ Broker: Select the MQ broker for which you want to view data.

Topics The number of topics found on the broker and listed in the table.

All Topics Table

Broker The name of the broker.

Topic String The name of the topic.

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 critical and warning alerts.

Publish Count The number of messages sent by publishers to the topic on the broker.

Subscription Count The number of messages received by subscribers from the topic on the broker.

Admin Topic Name The name of the topic object.

Cluster Name The cluster that is to be used for the propagation of publications and subscription to publish/subscribe cluster-connected queue managers for this topic.

Cluster Pub Route The routing behavior used for the topic.


Comm Info Name	The name of the communication information object.
Default Persistence	The setting for default persistence (either PERSISTENT or NOT PERSISTENT).
Default Priority	The resolved default priority of messages published to the topic.
Default Put Response Type	The default put response (synchronous or asynchronous).
Durable Subscriptions	Displays whether or not the applications are permitted to make durable subscriptions to the topic.
Inhibit Publications	Displays whether publications are allowed for the topic.
Inhibit Subscriptions	Displays whether subscriptions are allowed for the topic.
Model Durable Queue	The name of the model queue used for managed durable subscriptions.
Model Non Durable Queue	The name of the model queue used for managed non-durable subscriptions.
Multicast	Displays whether multicast is used for the topic.
NPM Delivery	Lists the delivery mechanism for non-persistent messages published to this topic (ALL, ALL_DUR, and ALL_AVAIL).
PM Delivery	Lists the delivery mechanism for persistent messages published to this topic (ALL, ALL_DUR, and ALL_AVAIL).
Publication Scope	Lists whether this queue manager propagates subscriptions for this topic to queue managers as part of hierarchy or as part of publish cluster.
Retained Publication	Lists whether there is a retained publication for this topic.
Subscription Scope	Lists whether this queue manager propagates subscriptions for this topic to queue managers as part of hierarchy or as part of subscribe cluster.
Use Dead letter Queue	Lists whether the dead-letter queue is used when publication messages cannot be delivered to their correct subscriber queue.
Expired	When checked, performance data has not been received in the time specified in the Duration region on the RTView Configuration Application > Solution Package Configuration > IBM MQ > DATA STORAGE tab.
Timestamp	The date and time of the last data update.

IBM MQ Topics Heatmap

This display allows you to view current alert status and performance metrics of all topics on a specific MQ broker in a heatmap format.

Use the **Metric** drop-down menu to view the heatmap based on **Alert Severity**, **Alert Count**, **Publish Count**, and **Subscription Count**.

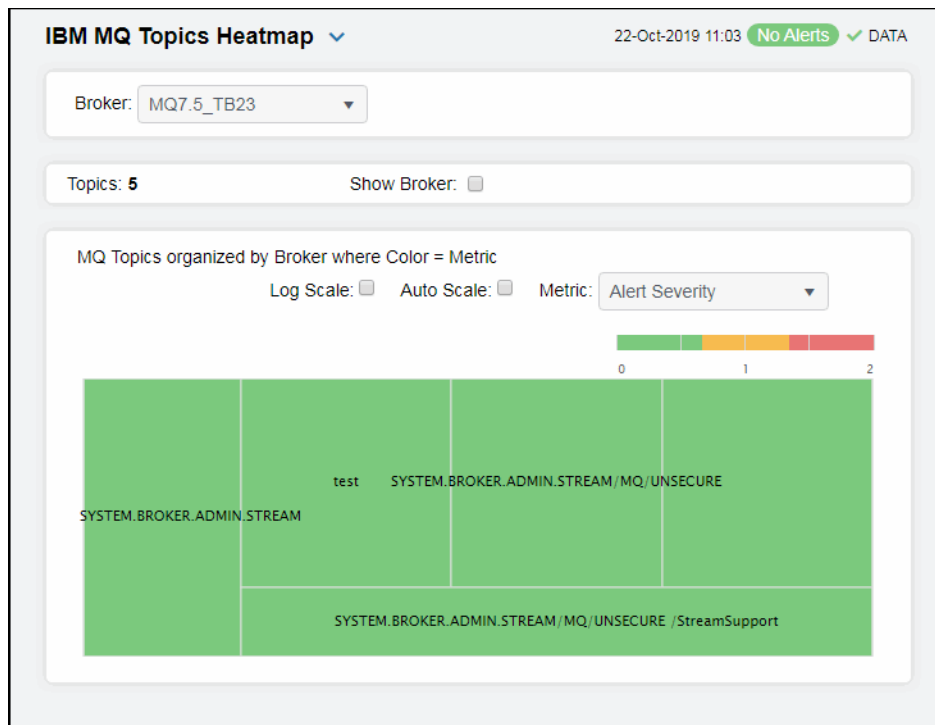
Each rectangle in the heatmap represents a different topic, where the rectangle color indicates the most critical alert state for items associated with that topic.

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.

Investigate a topic by clicking a rectangle which opens the “[IBM Topic Summary](#)” 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.



Filter By:

Broker: Select the MQ broker for which you want to view data.

Fields and Data

Topics Displays the number of topics found by the filter and listed in the heatmap.







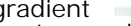
Show Broker Select this check box to display the names of the broker at the top of each rectangle in the heatmap.

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.

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 Scale** is not selected.

Metric	Choose a metric to view in the display. For details about the data, refer to vendor documentation.
Alert Severity	<p>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:</p> <ul style="list-style-type: none">  (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	<p>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.</p>
Publish Count	<p>The number of messages sent by publishers to the topic 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 number of published topics in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>
Subscription Count	<p>The number of messages received by subscribers from a topic 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 number of subscriptions in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>

IBM Topic Summary

This display allows you to view current alert status and performance metrics for a specific topic on a specific MQ broker. Select a broker and a topic string from the drop-down menus to view data for the selected topic string. Clicking on the information boxes at the top of the display (such as **Publish Count** and **Subscription Count**) opens the “[IBM MQ Topics Table](#)” display, where you can sort and compare the performance values of all topics.

The trend graph traces the **Metric Trends (Publish Count and Subscription Count)** for the selected topic. 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 can also 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 Topic Summary ▾
22-Oct-2019 13:55 3 Alerts ✓ DATA

Broker: vmrh5-1_legacy ▾

Topic String: SYSTEM.BROKER.ADMIN.STREAM/MQ/QM_0001 /StreamSupport ▾

Publish Count
0

Subscription Count
1

Metric Trends
Log Scale: 15 minutes ⌚

■ Publish Count ■ Subscription Count

Critical/Warning: 3/0			Connected: true	Expired: false
Admin Topic Name:	Cluster Name:	Cluster Pub Route:		
Comm Info Name:	Default Persistence: NOT PERSISTENT			
Default Priority: 0	Default Put Response Type: SYNC	Durable Subscriptions: DURABLE		
Inhibit Publications: ALLOWED	Inhibit Subscriptions: ALLOWED			
Model Durable Queue: SYSTEM.DURABLE.MODEL.QUEUE				
Model Non Durable Queue: SYSTEM.NDURABLE.MODEL.QUEUE			Multicast:	
NPM Delivery: ALL AVAIL	PM Delivery: ALL DURABLE	Publication Scope: ALL		
Retained Publication: YES	Subscription Scope: ALL	Use Dead Letter Queue:		

Last Update: **22-Oct-2019 13:55:39**

Filter By:

- Broker:** Select the MQ broker for which you want to view data.
- Topic String** Select the topic for which you want to view data.

Metric Trends Graph

Traces the following for the selected topic:

Publish Count: Traces the number of messages sent by publishers to the topic on the broker.


Subscription Count: Traces the number of messages received by subscribers from the topic 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.

Time Settings

Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

Critical/Warning

The total number of critical and warning alerts.

Connected

When checked, denotes that the broker is connected.

Expired

When checked, performance data has not been received in the time specified in the **Duration** region on the RTView Configuration Application > **Solution Package Configuration** > **IBM MQ** > **DATA STORAGE** tab.

Admin Topic Name

The name of the topic object.

Comm Info Name	The name of the communication information object.
Default Priority	The resolved default priority of messages published to the topic.
Durable Subscriptions	Displays whether or not the applications are permitted to make durable subscriptions to the topic.
Model Durable Queue	The name of the model queue used for managed durable subscriptions.
Model Non Durable Queue	The name of the model queue used for managed non-durable subscriptions.
NPM Delivery	Lists the delivery mechanism for non-persistent messages published to this topic (ALL, ALL_DUR, and ALL_AVAIL).
Retained Publication	Lists whether there is a retained publication for this topic.
Cluster Name	The cluster that is to be used for the propagation of publications and subscription to publish/subscribe cluster-connected queue managers for this topic.
Default Persistence	The setting for default persistence (either PERSISTENT or NOT PERSISTENT).
Default Put Response Type	The default put response (synchronous or asynchronous).
Inhibit Publications	Displays whether publications are allowed for the topic.
PM Delivery	Lists the delivery mechanism for persistent messages published to this topic (ALL, ALL_DUR, and ALL_AVAIL).
Subscription Scope	Lists whether this queue manager propagates subscriptions for this topic to queue managers as part of hierarchy or as part of subscribe cluster.
Cluster Pub Route	The routing behavior used for the topic.
Inhibit Subscriptions	Displays whether subscriptions are allowed for the topic.
Multicast	Displays whether multicast is used for the topic.
Publication Scope	Lists whether this queue manager propagates published messages for this topic to queue managers as part of hierarchy or as part of publish cluster.
Use Dead Letter Queue	Lists whether the dead-letter queue is used when publication messages cannot be delivered to their correct subscriber queue.
Last Update	The date and time of the last data update.

IBM MQ Topic Configs View

This View allows you to view performance and utilization metrics for all of your IBM MQ topic configs. Clicking **IBM MQ Topic Configs** in the left/navigation menu opens the “[IBM MQ Topic Configs Table](#)” display, which allows you to view current alert status and performance metrics of all topic configurations on a specific MQ broker in a table format.

- “[IBM MQ Topic Config Summary](#)”: Clicking **MQ Topic Config Summary** from the left/navigation menu opens this display, which allows you to view current alert status and performance metrics for a specific topic configuration on a specific MQ broker.

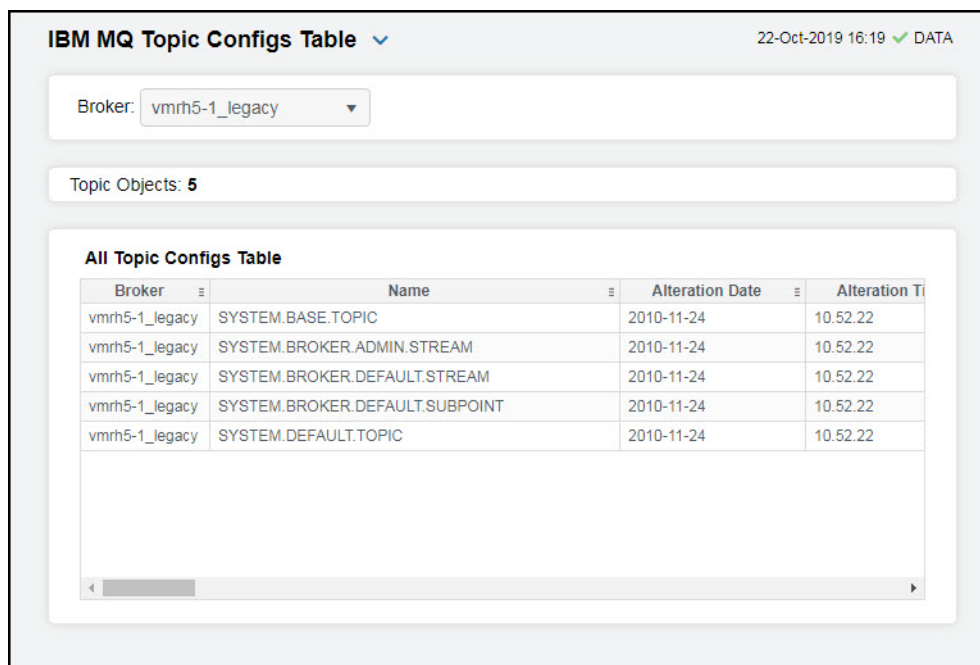
IBM MQ Topic Configs Table

This display allows you to view current alert status and performance metrics of all topic configurations on a specific MQ broker in a table format.

Each row in the table contains data for a particular topic config. 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**.

You can also 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 topic by double-clicking a row, which opens the “[IBM MQ Topic Config Summary](#)” display.



IBM MQ Topic Configs Table 22-Oct-2019 16:19 ✓ DATA

Broker:

Topic Objects: **5**

All Topic Configs Table

Broker	Name	Alteration Date	Alteration T
vmrh5-1_legacy	SYSTEM.BASE.TOPIC	2010-11-24	10.52.22
vmrh5-1_legacy	SYSTEM.BROKER.ADMIN.STREAM	2010-11-24	10.52.22
vmrh5-1_legacy	SYSTEM.BROKER.DEFAULT.STREAM	2010-11-24	10.52.22
vmrh5-1_legacy	SYSTEM.BROKER.DEFAULT.SUBPOINT	2010-11-24	10.52.22
vmrh5-1_legacy	SYSTEM.DEFAULT.TOPIC	2010-11-24	10.52.22

Filter By:

Broker: Select the MQ broker for which you want to view data.

Topic Objects	The number of topics found on the broker and listed in the table.
Fields and Data	
Broker	The name of the broker.
Name	The name of the topic.
Alteration Date	The date on which the information was last altered.
Alteration Time	The time at which the information was last altered.
Cluster Name	The cluster that is to be used for the propagation of publications and subscription to publish/subscribe cluster-connected queue managers for this topic.
Cluster Object State	The current state of the clustered topic definition.
Cluster Publish Route	The routing behavior used for the topic.
Comm Info Name	The name of the communication information object.
Custom	The custom attribute for new features.
Default Persistence	The setting for default persistence (either PERSISTENT or NOT PERSISTENT)
Default Priority	The resolved default priority of messages published to the topic.
Default Put Response Type	The default put response (synchronous or asynchronous).
Description	Description of the topic object.
Durable Sub	Displays whether or not the applications are permitted to make durable subscriptions to the topic.
Inhibit Pub	Displays whether publications are allowed for the topic.
Inhibit Sub	Displays whether subscriptions are allowed for the topic.
Model Durable Queue	The name of the model queue used for managed durable subscriptions.
Model Non Durable Queue	The name of the model queue used for managed non-durable subscriptions.
Multicast	Displays whether multicast is used for the topic.
NPM Delivery	Lists the delivery mechanism for non-persistent messages published to this topic (ALL, ALL_DUR, and ALL_AVAIL).
PM Delivery	Lists the delivery mechanism for persistent messages published to this topic (ALL, ALL_DUR, and ALL_AVAIL).
Proxy Sub	Lists whether a proxy subscription is to be sent for this topic, even if no local subscriptions exist.
Pub Scope	Lists whether this queue manager propagates published messages for this topic to queue managers as part of hierarchy or as part of publish cluster.
Sub Scope	Lists whether this queue manager propagates subscriptions for this topic to queue managers as part of hierarchy or as part of subscribe cluster.

Topic Type	The type of topic.
Topic String	The name of the topic string.
Use Dead Letter Queue	Lists whether the dead-letter queue is used when publication messages cannot be delivered to their correct subscriber queue.
Wildcard Operation	Lists how subscriptions made using wildcard topic names that are less specific than the topic string at this topic object are handled (PASSTHRU or BLOCK).
Expired	When checked, performance data has not been received in the time specified in the Duration region on the RTView Configuration Application > Solution Package Configuration > IBM MQ > DATA STORAGE tab.
Timestamp	The date and time of the last data update.

IBM MQ Topic Config Summary

This display allows you to view current alert status and performance metrics for a specific topic configuration on a specific MQ broker. Select a broker and the name of the topic config from the drop-down menus to view data for the selected topic config. Clicking on the information boxes at the top of the display (such as **Topic String** and **Description**) opens the “[IBM MQ Topic Configs Table](#)” display, where you can sort and compare the performance values of all topics.

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.

Filter By:

- Broker:** Select the MQ broker for which you want to view data.
- Name** Select the topic for which you would like to view data.

Fields and Data

Topic String	The name of the topic string for the selected topic object.
Description	The description of the topic object.
Alteration Date	The date on which the information was last altered.
Alteration Time	The time at which the information was last altered.
Critical/Warning	The total number of critical and warning alerts.
Connected	When checked, denotes that the broker is connected.
Expired	When checked, performance data has not been received in the time specified in the Duration region on the RTView Configuration Application > Solution Package Configuration > IBM MQ > DATA STORAGE tab.
Cluster Name	The cluster that is to be used for the propagation of publications and subscription to publish/subscribe cluster-connected queue managers for this topic.
Comm Info Name	The name of the communication information object.
Default Priority	The resolved default priority of messages published to the topic.
Inhibit Pub	Displays whether publications are allowed for the topic.
Model Non Durable Queue	The name of the model queue used for managed non-durable subscriptions.
PM Delivery	Lists the delivery mechanism for persistent messages published to this topic (ALL, ALL_DUR, and ALL_AVAIL).
Sub Scope	Lists whether this queue manager propagates subscriptions for this topic to queue managers as part of hierarchy or as part of subscribe cluster.
Wildcard Operation	Lists how subscriptions made using wildcard topic names that are less specific than the topic string at this topic object are handled (PASSTHRU or BLOCK).
Cluster Object State	The current state of the clustered topic definition.
Custom	The custom attribute for new features.
Default Put Response Type	The default put response (synchronous or asynchronous).
Inhibit Sub	Displays whether subscriptions are allowed for the topic.
Multicast	Displays whether multicast is used for the topic.
Proxy Sub	Lists whether a proxy subscription is to be sent for this topic, even if no local subscriptions exist.
Topic Type	The type of topic.
Cluster Publish Route	The routing behavior used for the topic.
Default Persistence	The setting for default persistence (either PERSISTENT or NOT PERSISTENT)
Durable Sub	Displays whether or not the applications are permitted to make durable subscriptions to the topic.
Model Durable Queue	The name of the model queue used for managed durable subscriptions.
NPM Delivery	Lists the delivery mechanism for non-persistent messages published to this topic (ALL, ALL_DUR, and ALL_AVAIL).

Pub Scope	Lists whether this queue manager propagates published messages for this topic to queue managers as part of hierarchy or as part of publish cluster.
Use Dead Letter Queue	Lists whether the dead-letter queue is used when publication messages cannot be delivered to their correct subscriber queue.
Last Update	The date and time of the last data update.

IBM MQ Topic Publishers View

The displays in this View present performance and utilization metrics for your Topic Publishers. Clicking **IBM MQ Topic Publishers** in the left/navigation menu opens the [“IBM MQ Topic Publishers Table”](#) display, which allows you to view current alert status and performance metrics of all topic publishers on a specific MQ broker in a table format. The other available displays in this View are:

- [“IBM MQ Topic Publishers Heatmap”](#): Clicking **MQ Topic Publishers Heatmap** opens this display, which allows you to view current alert status and performance metrics of all topic publishers on a specific MQ broker in a heatmap format.
- [“IBM MQ Topic Publisher Summary”](#): Clicking **MQ Topic Publisher Summary** opens this display, which allows you to view current alert status and performance metrics for a specific topic publisher on a specific MQ broker.

IBM MQ Topic Publishers Table

This display allows you to view current alert status and performance metrics of all topic publishers on a specific MQ broker in a table format.

Each row in the table contains data for a particular publisher. 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**.

You can also 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 topic by double-clicking a row, which opens the [“IBM MQ Topic Publisher Summary”](#) display.

IBM MQ Topic Publishers Table 23-Oct-2019 10:12 No Alerts DATA

Broker: - All - Topic String: - All -

Topic Publishers: 1

All Topic Publishers Table

Broker	Topic String	
MQ9.1_TB29	dev	414d5

Filter By:

Broker: Select the MQ broker for which you want to view data.

Topic String Select the topic for which you want to view data.

Topic Publishers The number of topic publishers found on the broker and listed in the table.

All Topic Publishers Table

Broker The name of the broker.

Topic String The name of the topic string.

Active Connection The currently active ConnectionId (CONNID) associated with the handle that has this topic open for publish.

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 critical and warning alerts.

Delta Publish Count The current number of applications publishing to the topic on the broker since the last data update.

Delta Timestamp The difference in time (in milliseconds) between the most recent update (Last Update/ Timestamp) and the previous update.

Last Publication Date The date on which this publisher last sent a message.

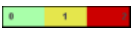
Last Publication Time	The time at which this publisher last sent a message.
Multicast Reliability Indicator	The multicast reliability indicator.
Publish Count	The number of applications publishing to the topic on the broker.
Rate Publish Count	The rate of messages sent by publishers to the topic on the broker.
Expired	When checked, performance data has not been received in the time specified in the Duration region on the RTView Configuration Application > Solution Package Configuration > IBM MQ > DATA STORAGE tab.
Timestamp	The date and time of the last data update.




IBM MQ Topic Publishers Heatmap

View current alert status and performance metrics of all topic publishers on a specific broker.

Use the **Metric** drop-down menu to view the heatmap based on **Alert Severity**, **Alert Count**, and **Rate Publish Count**.

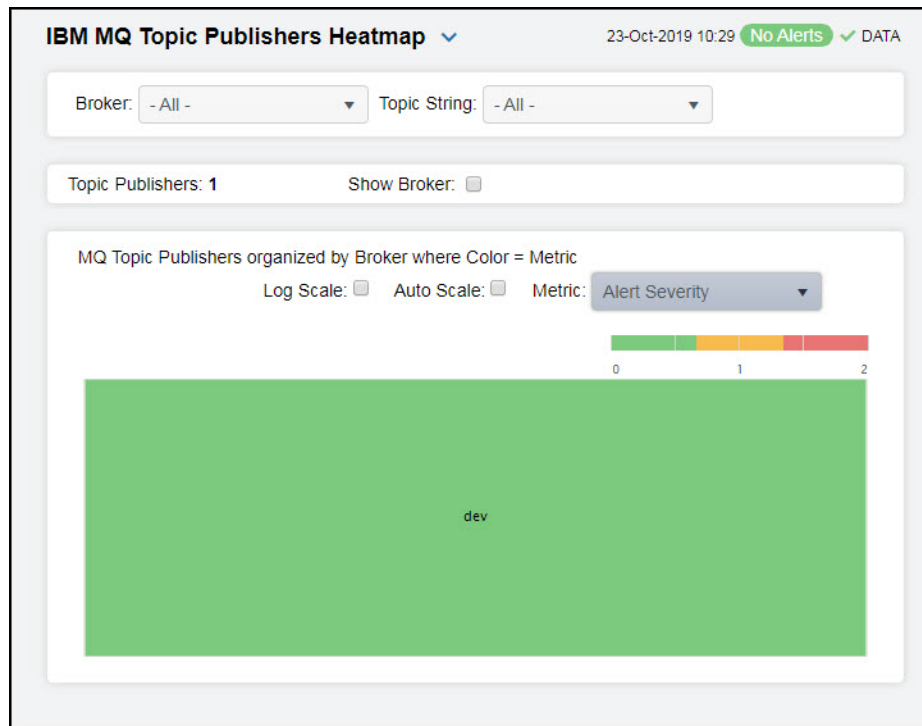
Each rectangle in the heatmap represents a different topic publisher, where the rectangle color indicates the most critical alert state for items associated with that topic publisher.

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.

Investigate a topic by clicking a rectangle which opens the ["IBM MQ Topic Publisher Summary"](#) display.

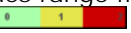




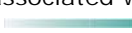
Mouse-over rectangles to view more details about host performance and status. Toggle between the commonly accessed Table and Heatmap displays by clicking the drop down list on the display title.

**Filter By:**

- Broker** Select the MQ broker containing the topic string for which you want to view data or select **All** to view all available topic strings for all brokers.
- Topic String** Select the topic string for which you want to view data or select **All** to view all topic strings.

Fields and Data

- Topic Publishers** The number of topic publishers found as a result of the filter and listed in the heatmap.
- Show Broker** Select this check box to display the names of the brokers 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.
- 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 Scale** is not selected.
- Metric** Choose a metric to view in the display. For details about the data, refer to vendor documentation.

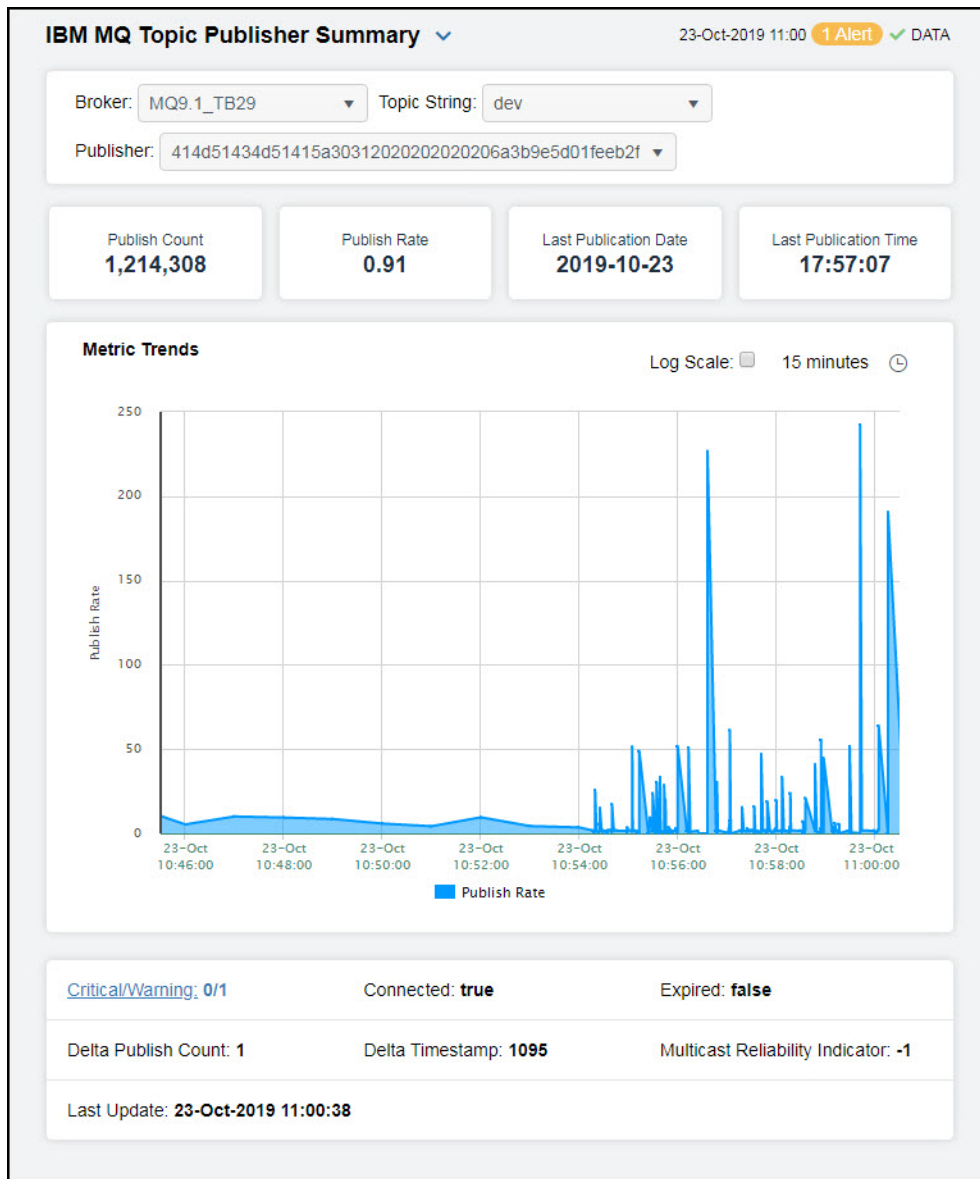
Alert Severity	<p>The current alert severity for items associated with the rectangle. Values range from 0 - 2, as indicated in the color gradient  bar, where 2 is the highest Alert Severity:</p> <ul style="list-style-type: none">  (2) Red indicates that one or more metrics exceeded their ALARM LEVEL 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	<p>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.</p>
Rate Publish Count	<p>The rate of applications publishing to the topics 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 rate of published topics in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>

IBM MQ Topic Publisher Summary

View current alert status and performance metrics for a specific topic publisher on a specific broker. Select a broker, a topic string, and a publisher from the drop-down menus to view data for the selected topic publisher. Clicking on the information boxes at the top of the display (such as **Publish Count** and **Publish Rate**) opens the ["IBM MQ Topic Publishers Table"](#) display, where you can sort and compare the performance values of all topic publishers.

The trend graph traces the **Metric Trends (Publish Rate)** for the selected topic publisher. 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. You can hover over the trend graph to see the values at a particular time.

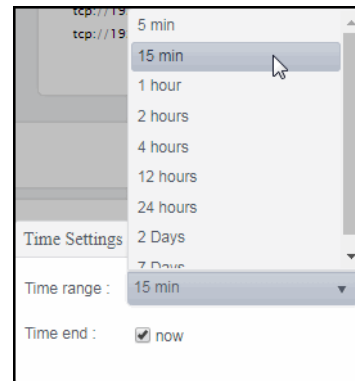
**Filter By:**

- Broker:** Select the MQ broker containing the topic string and publisher for which you want to view data.
- Topic String** Select the topic string containing the publisher for which you want to view data.
- Publisher** Select the publisher for which you want to view data.

Fields and Data

- Publish Count** The number of messages sent by publishers to the topic on the broker.
- Publish Rate** The rate of messages sent by publishers to the topic on the broker.

Last Publication Date	The date on which this publisher last sent a message.
Last Publication Time	The time at which this publisher last sent a message.
Metric Trends Graph	Traces the following for the selected topic: Publish Rate: Traces the rate of messages sent by publishers to the topic 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.
Time Settings	Select a time range from the drop down menu varying from 5 Minutes to Last 7 Days . By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

Critical/ Warning	The total number of critical and warning alerts.
Connected	When checked, denotes that the broker is connected.
Expired	When checked, performance data has not been received in the time specified in the Duration region on the RTView Configuration Application > Solution Package Configuration > IBM MQ > DATA STORAGE tab.

Delta Publish Count	The current number of applications publishing to the topic on the broker since the last data update.
Delta Timestamp	The difference in time (in milliseconds) between the most recent update (Last Update/ Timestamp) and the previous update.
Multicast Reliability Indicator	The multicast reliability indicator.
Last Update	The date and time of the last data update.

IBM MQ Topic Subscribers View

This View allows you to view performance and utilization metrics for all of your IBM MQ topic subscribers. Clicking **IBM MQ Topic Subscribers** in the left/navigation menu opens the [“IBM MQ Topic Subscribers Table”](#) display, which allows you to view current alert status and performance metrics of all topic subscribers on a specific MQ broker in a table format.

- [“IBM MQ Topic Subscribers Heatmap”](#): Clicking **MQ Topic Subscribers Heatmap** in the left/navigation menu opens this display, which allows you to view current alert status and performance metrics of all topic subscribers on a specific MQ broker in a heatmap format.
- [“IBM MQ Topic Subscriber Summary”](#): Clicking **IBM MQ Topic Subscriber Summary** in the left/navigation menu opens this display, which allows you to view current alert status and performance metrics for a specific topic subscriber on a specific MQ broker.

IBM MQ Topic Subscribers Table

This display allows you to view current alert status and performance metrics of all topic subscribers on a specific MQ broker in a table format.

Each row in the table contains data for a particular subscriber. 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**.

You can also 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 topic by double-clicking a row, which opens the [“IBM MQ Topic Subscriber Summary”](#) display.

IBM MQ Topic Subscribers Table 23-Oct-2019 14:39 2 Alerts DATA

Broker: MQ9.1_TB29 Topic String: dev

Topic Subscribers: 2

All Topic Subscribers Table

Broker	Topic String	Subscription ID
MQ9.1_TB29	dev	414d5
MQ9.1_TB29	dev	414d5

Filter By:

Broker: Select the broker containing the topic string for which you want to view data or select **All** to include all topic strings.

Topic String Select the topic string for which you want to view data or select **All** to view all topic strings.

Topic Subscribers The number of topic publishers found using the filter and listed in the table.

All Topic Subscribers Table

Broker The name of the broker.

Topic String The name of the topic string.

Subscription ID The ID of the subscription.

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 critical and warning alerts.

Active Connection The ConnId of the HConn that currently has this subscription open.


Delta Message Count The number of messages since the last data update.




Delta Time Stamp	The difference in time (in milliseconds) between the most recent update (Last Update/ Timestamp) and the previous update.
Durable	When NO, the subscription is removed when the application that created it is closed or disconnected from the queue manager. When YES, the subscription persists even when the creating application is no longer running or has been disconnected. The subscription is reinstated when the queue manager is restarted.
Last Message Date	The date that a message was last sent to the destination specified by the subscription.
Last Message Time	The time when a message was last sent to the destination specified by the subscription.
Message Count	The number of messages put to the destination specified by this subscription.
Multicast Reliability Indicator	The multicast reliability indicator.
Rate Message Count	The rate of messages put to the destination specified by this subscription.
Resume Date	The date of the most recent MQSUB API call that connected to this subscription.
Resume Time	The time of the most recent MQSUB API call that connected to this subscription.
Subscription Type	The type of subscription, which indicates how the subscription was created.
Subscription User ID	The user ID that owns the subscription.
Expired	When checked, performance data has not been received in the time specified in the Duration region on the RTView Configuration Application > Solution Package Configuration > IBM MQ > DATA STORAGE tab.
Timestamp	The date and time of the last data update.

IBM MQ Topic Subscribers Heatmap

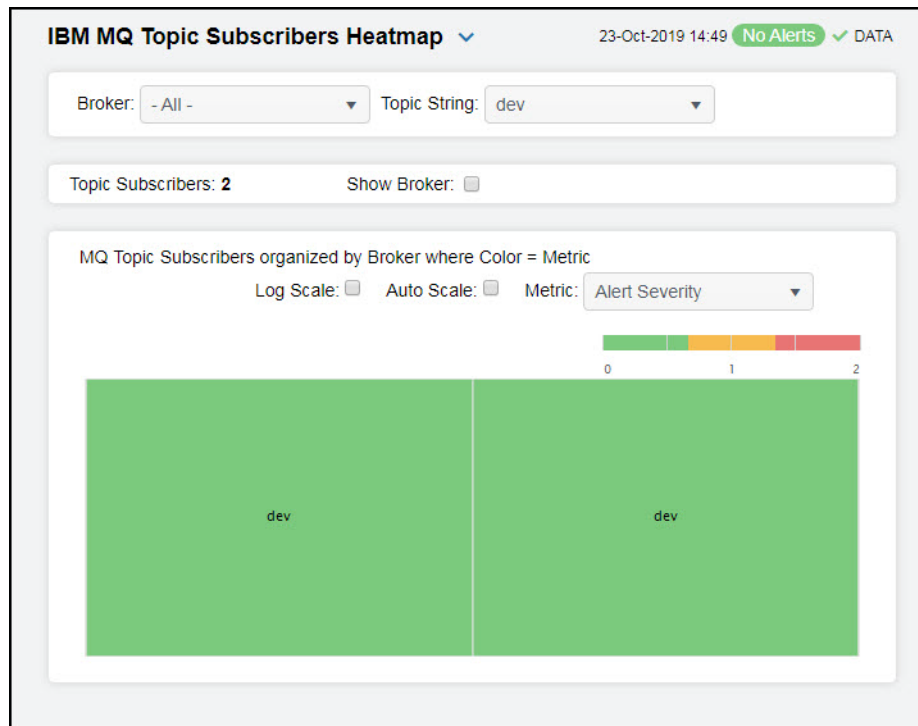
View current alert status and performance metrics of all topic subscribers on a specific MQ broker.

Each rectangle in the heatmap represents a different topic subscriber, where the rectangle color indicates the most critical alert state for items associated with that topic subscriber, and the rectangle size represents the tablespace allocation size for the topic subscriber.

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.

Use the **Show Broker** 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 ["IBM MQ Topic Subscriber Summary"](#) display.

**Filter By:**

Broker: Select the MQ broker for which you want to view data..

Fields and Data

Topic Subscribers The number of topic subscribers found by the filter and listed in the heatmap.





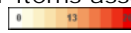

Show Broker Select this check box to display the names of the brokers at the top of each rectangle in the heatmap.

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.

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 Scale** is not selected.

Metric Choose a metric to view in the display. For details about the data, refer to vendor documentation.

Alert Severity	<p>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:</p> <ul style="list-style-type: none">  (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	<p>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.</p>
Rate Message Count	<p>The rate of messages put to the destination specified by the subscription. The color 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 published topics in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>

IBM MQ Topic Subscriber Summary

Use this display to view current alert status and performance metrics for a specific topic subscriber on a specific MQ broker. Select a broker, a topic string, and a subscription ID from the drop-down menus to view data for the selected topic subscriber. Clicking on the information boxes at the top of the display (such as **Message Count** and **Message Rate**) opens the ["IBM MQ Topic Subscribers Table"](#) display, where you can sort and compare the performance values of all topic subscribers.

The trend graph traces the **Metric Trends (Message Rate)** for the selected topic subscriber. 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. You can hover over the trend graph to see the values at a particular time.

IBM MQ Topic Subscriber Summary ▾
23-Oct-2019 15:52 1 Alert ✓ DATA

Broker: MQ9.1_TB29 ▾
Topic String: dev ▾

SubscriptionID: 414d51204d51415a30312020202020205c718a5d040b4b22 ▾

Message Count
3,606,924

Message Rate
0.00

Last Message Date
2019-10-23

Last Message Time
22:49:28

Resume Date
2019-10-15

Resume Time
17:42:17

Metric Trends
Log Scale: 15 minutes 🕒

[Critical/Warning: 0/1](#)
Connected: **true**
Expired: **false**

Active Connection: **414d51434d51415a303120202020206a3b9e5d01fdeb2f**

Delta Message Count: **0**
Delta Time Stamp: **2063**
Durable: **YES**

Multicast Reliability Indicator: **-1**
Subscription Type: **API**
Subscription User ID:

Last Update: **23-Oct-2019 15:53:01**

Filter By:

Broker: Select the MQ broker for which you want to view data.

Topic String Select the topic for which you want to view data.

SubscriptionID Select the subscription ID for which you want to view data.

Fields and Data

Connected When checked, denotes that the broker is connected.

Expired When checked, performance data has not been received in the time specified in the **Duration** region on the RTView Configuration Application > **Solution Package Configuration** > **IBM MQ** > **DATA STORAGE** tab.

Last Update The date and time of the last data update.

Fields and Data

Message Count The number of messages put to the destination that are specified by this subscription.

Message Rate The rate of messages put to the destination that are specified by this subscription.

Last Message Date The date that a message was last sent to the destination specified by the subscription.

Last Message Time The time when a message was last sent to the destination specified by the subscription.

Metric Trends Graph

Traces the following for the selected topic:


Message Rate: Traces the rate of messages put to the destination specified by this subscription.

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

Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

Critical/Warning	The number of critical and warning alerts.
Connected	When checked, denotes that the broker is connected.
Expired	When checked, performance data has not been received in the time specified in the Duration region on the RTView Configuration Application > Solution Package Configuration > IBM MQ > DATA STORAGE tab.
Active Connection	The ConnId of the HConn that currently has this subscription open.
Delta Message Count	The number of messages since the last data update.
Multicast Reliability Indicator	The multicast reliability indicator.
Delta Time Stamp	The difference in time (in milliseconds) between the most recent update (Last Update/ Timestamp) and the previous update.
Subscription Type	The type of subscription, which indicates how the subscription was created.
Durable	When NO , the subscription is removed when the application that created it is closed or disconnected from the queue manager. When YES , the subscription persists even when the creating application is no longer running or has been disconnected. The subscription is reinstated when the queue manager is restarted.

Subscription User ID	The user ID that owns the subscription.
Last Update	The date and time of the last data update.

IBM MQ Topic Publisher Totals View

This View allows you to view performance and utilization metrics for all of your IBM MQ topic publisher totals. Clicking **IBM MQ Topic Publishers Totals** in the left/navigation menu opens the “[IBM MQ Topic Publisher Totals Table](#)” display, which allows you to view current alert status and performance metrics based on the total messages sent by all publishers for each topic on a specific MQ broker in a table format.

- “[IBM MQ Topic Publisher Totals Heatmap](#)”: Clicking **MQ Topic Publisher Totals Heatmap** in the left/navigation menu opens this display, which allows you to view current alert status and performance metrics based on the total messages sent by all publishers for each topic on a specific MQ broker in a heatmap format.
- “[IBM MQ Topic Publisher Totals Summary](#)”: Clicking **MQ Topic Publisher Totals Summary** in the left/navigation menu opens this display, which allows you to view current alert status and performance metrics based on the total messages sent by a specific publisher on a specific MQ broker.

IBM MQ Topic Publisher Totals Table

This display allows you to view current alert status and performance metrics based on the total messages sent by all publishers for each topic on a specific MQ broker in a table format.

Each row in the table contains data for a particular topic publisher’s totals. 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**.

You can also 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 topic by double-clicking a row, which opens the “[IBM MQ Topic Publisher Totals Summary](#)” display.

IBM MQ Topic Publisher Totals Table 24-Oct-2019 10:29 No Alerts ✓ DATA

Broker: MQ9.1_TB29

Topic Publisher Totals: 1

All Topic Publishers Table

Broker	Topic String	Alert Level
MQ9.1_TB29	dev	✓

Filter By:

Broker: Select the MQ broker for which you want to view data.

Topic Publisher Totals The number of topic publishers found on the broker and listed in the table.

All Topic Publisher Totals Table

Broker The name of the broker.

Topic String The name of the topic string.

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 critical and warning alerts.

Delta Publish Count The current number of messages sent by publishers to the topic on the broker since the last data update.

Rate Publish Count The rate of messages sent by publishers to the topic on the broker.


Expired When checked, performance data has not been received in the time specified in the **Duration** region on the RTView Configuration Application > **Solution Package Configuration** > **IBM MQ** > **DATA STORAGE** tab.




Timestamp The date and time of the last data update.

IBM MQ Topic Publisher Totals Heatmap

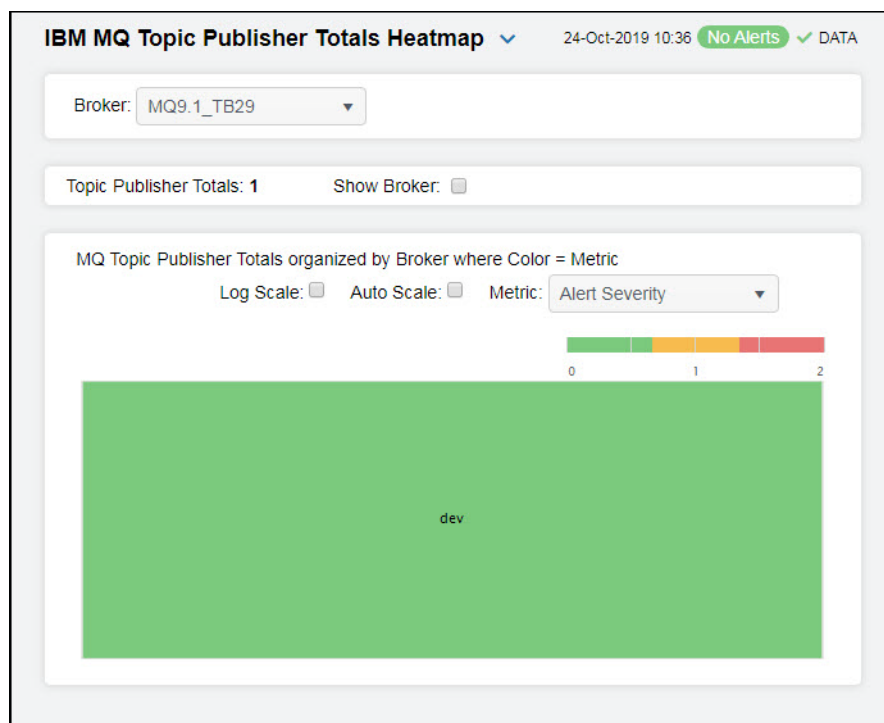
This display allows you to view current alert status and performance metrics based on the total messages sent by all publishers for each topic on a specific MQ broker in a heatmap format.

Each rectangle in the heatmap represents a different topic publisher totals, where the rectangle color indicates the most critical alert state for items associated with that topic publisher totals, and the rectangle size represents the tablespace allocation size for the topic publisher totals.

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.

Use the **Show Broker** 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 ["IBM MQ Topic Publisher Totals Summary"](#) display.



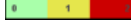





Filter By:

Broker: Select the broker for which you want to view data or select **All** to view data for all topic publishers for all brokers.

Fields and Data

Topic Publisher Totals	The number of topic publisher totals found on the broker and listed in the heatmap.
Show Broker	Select this check box to display the names of the broker at the top of each rectangle in the heatmap.

Heatmap

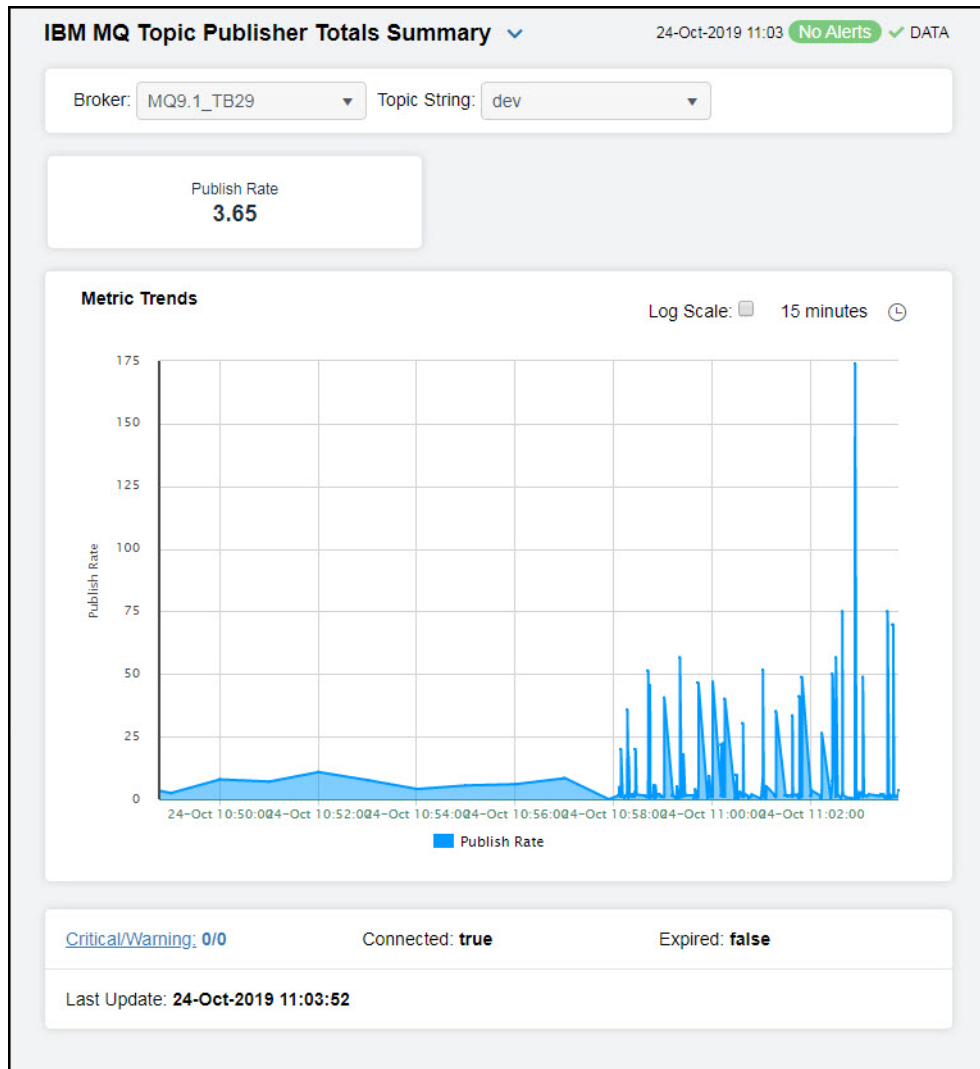
Log Scale	Select this check box to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not 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.
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 Scale is not selected.
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: <ul style="list-style-type: none">  (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.
Rate Publish Count	The rate of messages sent by publishers for the particular topic. The color 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 published topics in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

IBM MQ Topic Publisher Totals Summary

This display allows you to view current alert status and performance metrics based on the total messages sent by a specific publisher on a specific MQ broker. Select a broker and a topic string from the drop-down menus to view data for the selected topic publisher. Clicking on the **Publish Rate** information box at the top of the display opens the ["IBM MQ Topic Publisher Totals Table"](#) display, where you can sort and compare the performance values of all topic publisher totals.

The trend graph traces the **Metric Trends (Publish Rate)** for the selected topic publisher totals. 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. You can hover over the trend graph to see the values at a particular time.



Filter By:

Broker: Select the MQ broker for which you want to view data.

Topic String Select the topic string for which you want to view data.

Fields and Data

Publish Rate The rate of messages sent by the publisher.

Metric Trends Graph

Traces the following for the selected topic:

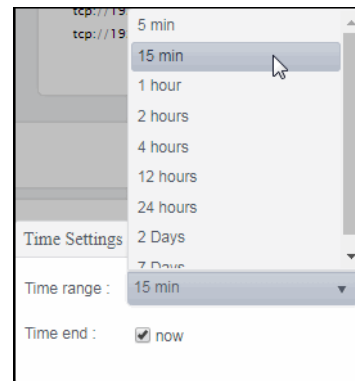
Publish Rate: Traces the rate of messages sent by the publisher.

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

Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

Critical/Warning The total number of critical and warning alerts.

Connected When checked, denotes that the broker is connected.

Expired When true, performance data has not been received in the time specified in the **Duration** region on the RTView Configuration Application > **Solution Package Configuration** > **IBM MQ** > **DATA STORAGE** tab.

Last Update The date and time of the last data update.

IBM MQ Topic Subscriber Totals View

This View contains displays that allow you to view performance and utilization metrics for all of your IBM MQ topic subscriber totals. Clicking **IBM MQ Topic Subscribers Totals** in the left/navigation menu opens the [“IBM MQ Topic Subscriber Totals Table”](#) display, which allows you to view current alert status and performance metrics based on the total messages received by all subscribers for each topic on a specific MQ broker in a table format.

- [“IBM MQ Topic Subscriber Totals Heatmap”](#): Clicking **IBM MQ Topic Subscribers Totals Heatmap** in the left/navigation menu opens this display, which allows you to view current alert status and performance metrics based on the total messages received by all subscribers for each topic on a specific MQ broker in a heatmap format.
- [“IBM Topic Subscriber Totals Summary”](#): Clicking **IBM MQ Topic Subscribers Totals Summary** in the left/navigation menu opens this display, which allows you to view current alert status and performance metrics based on the total messages received by all subscribers for a particular topic on a specific MQ broker.

IBM MQ Topic Subscriber Totals Table

This display allows you to view current alert status and performance metrics based on the total messages received by all subscribers for each topic on a specific MQ broker in a table format.

Each row in the table contains data for a particular topic subscriber's totals. 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**.

You can also 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 topic subscriber totals by double-clicking a row, which opens the [“IBM Topic Subscriber Totals Summary”](#) display.

IBM MQ Topic Subscriber Totals Table 24-Oct-2019 11:46 No Alerts DATA

Broker: MQ9.1_TB29

Topic Subscriber Totals: **2**

All Topic Subscribers Table

Broker	Topic String	Alert Level
MQ9.1_TB29	dev	✓
MQ9.1_TB29	SYSTEM.BROKER.ADMIN.STREAM/MQ/MQAZ01 /StreamSupport	✓

Filter By:

Broker: Select the broker for which you want to view data or select **All** to view subscriber totals for all brokers.

Topic Subscriber Totals The number of topic subscriber totals found on the broker and listed in the table.

All Topic Subscriber Totals Table

Broker The name of the broker.

Topic String The name of the topic string.

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 critical and warning alerts.

Delta Message Count The number of messages since the last data update.

Rate Message Count The rate of messages put to the destination specified by this subscription.


Expired When checked, performance data has not been received in the time specified in the **Duration** region on the RTView Configuration Application > **Solution Package Configuration** > **IBM MQ** > **DATA STORAGE** tab.




Timestamp The date and time of the last data update.

IBM MQ Topic Subscriber Totals Heatmap

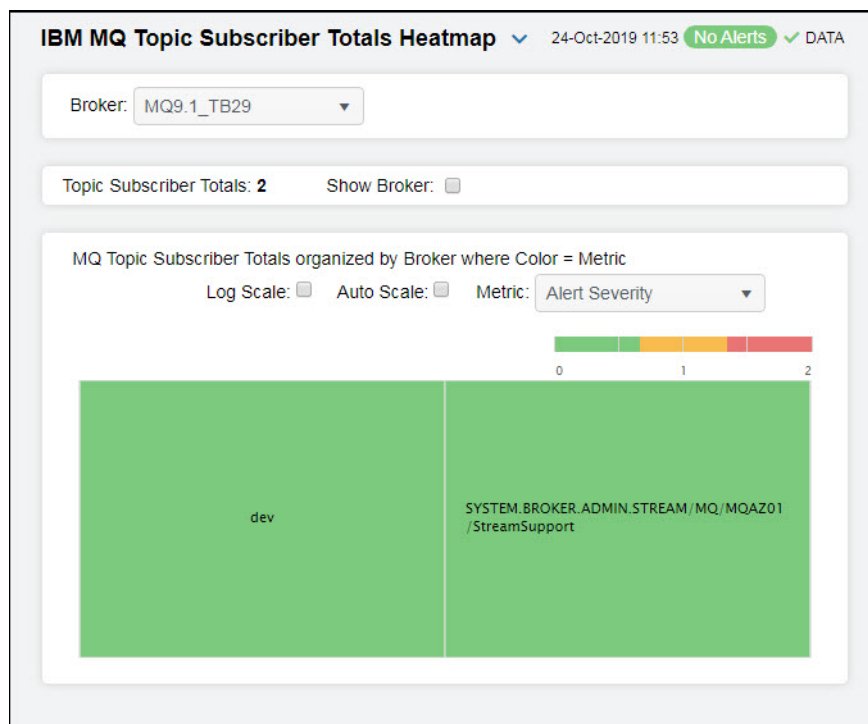
This display allows you to view current alert status and performance metrics based on the total messages received by all subscribers for each topic on a specific MQ broker in a heatmap format.

Each rectangle in the heatmap represents a different topic subscriber totals, where the rectangle color indicates the most critical alert state for items associated with that topic subscriber totals, and the rectangle size represents the tablespace allocation size for the topic subscriber totals.

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.

Use the **Show Broker** 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 ["IBM Topic Subscriber Totals Summary"](#) display.









Filter By:

Broker

Select the broker for which you want to view data or select All to view all subscriber totals on all brokers.

Fields and Data

Topic Subscriber Totals	The number of topic subscriber totals found on the broker and listed in the heatmap.
Show Broker	Select this check box to display the name of the broker 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.
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 Scale is not selected.
Metric	Choose a metric to view in the display. For details about the data, refer to vendor documentation.
Alert Severity	<p>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:</p> <ul style="list-style-type: none">  (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.
Rate Message Count	The rate of messages put to the destination specified by the subscription. The color 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 published topics in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

IBM Topic Subscriber Totals Summary

This display allows you to view current alert status and performance metrics based on the total messages received by all subscribers for a particular topic on a specific MQ broker. Select a broker and a topic string from the drop-down menus to view data for the selected topic subscriber. Clicking on the **Message Rate** information box at the top of the display opens the ["IBM MQ Topic Subscriber Totals Table"](#) display, where you can sort and compare the performance values of all topic subscriber totals.

The trend graph traces the **Metric Trends (Message Rate)** for the selected topic subscriber totals. 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. You can hover over the trend graph to see the values at a particular time.

Filter By:

Broker: Select the MQ broker for which you want to view data.

Topic String Select the topic string for which you want to view data.

Fields and Data

Message Rate The rate of messages put to the destination specified by this subscription.

Metric Trends Graph

Traces the following for the selected topic:


Message Rate: Traces the rate of messages put to the destination specified by this subscription.

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

Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

Critical/Warning	The total number of critical and warning alerts.
Connected	When checked, denotes that the broker is connected.
Expired	When true, performance data has not been received in the time specified in the Duration region on the RTView Configuration Application > Solution Package Configuration > IBM MQ > DATA STORAGE tab.
Last Update	The date and time of the last data update.

IBM MQ Subscriptions View

This View contains displays that allow you to view performance and utilization metrics for all of your IBM MQ subscriptions. Clicking **IBM MQ Subscriptions** in the left/navigation menu opens the ["IBM MQ Subscriptions Table"](#) display, which allows you to view current alert status and performance metrics based on the total subscriptions on a specific MQ broker in a table format.

- ["IBM MQ Subscriptions Heatmap"](#): Clicking **MQ Subscriptions Heatmap** in the left/navigation menu opens this display, which allows you to view current alert status and performance metrics based on the total subscriptions on a specific MQ broker in a heatmap format.
- ["IBM MQ Subscription Summary"](#): Clicking **MQ Subscriptions Summary** in the left/navigation menu opens this display, which allows you to view current alert status and performance metrics for a specific subscription on a specific MQ broker.

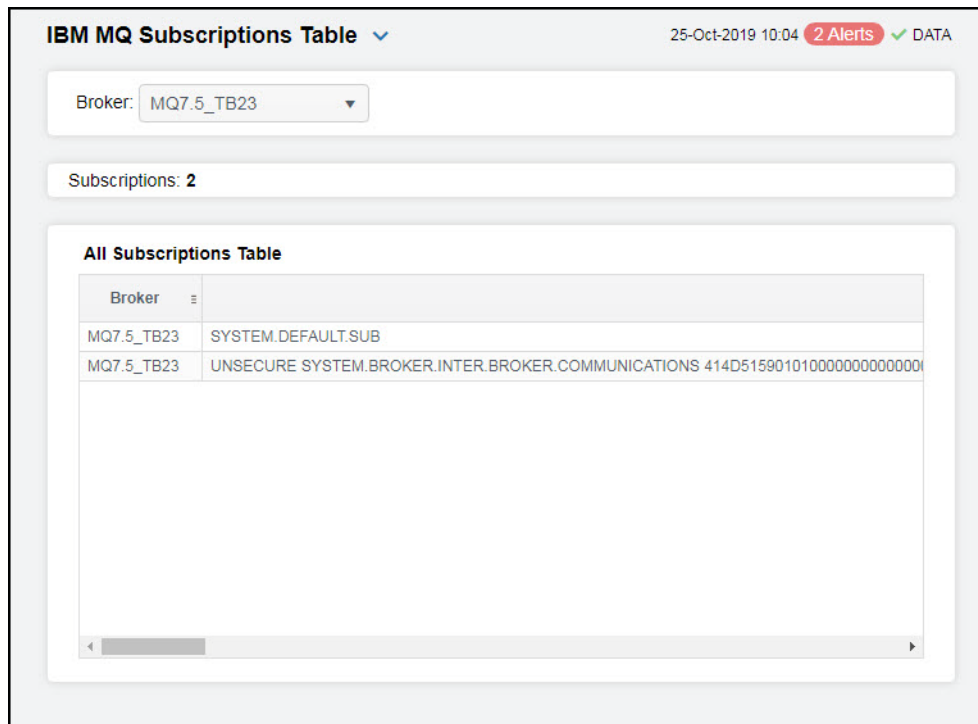
IBM MQ Subscriptions Table

This display allows you to view current alert status and performance metrics based on the total subscriptions on a specific MQ broker in a table format.

Each row in the table contains data for a particular subscription. 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**.

You can also 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 subscription by double-clicking a row, which opens the ["IBM MQ Subscription Summary"](#) display.



Filter By:

Broker: Select the broker for which you want to view data or select **All** to view subscriptions for all brokers.

Subscriptions The number of subscriptions found on the broker and listed in the table.

All Subscriptions Table

Broker The name of the broker.

Name The name of the subscription.

Subscription ID The ID of the subscription.

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 critical and warning alerts.

Active Connection The currently active ConnectionId (CONNID) associated with the handle that has this topic open for publish.

Delta Message Count The number of messages since the last data update.


Deltatime_stamp The difference in time (in milliseconds) between the most recent update (Last Update/ Timestamp) and the previous update.




Durable	When NO, the subscription is removed when the application that created it is closed or disconnected from the queue manager. When YES, the subscription persists even when the creating application is no longer running or has been disconnected. The subscription is reinstated when the queue manager is restarted.
Last Message Date	The date that a message was last sent to the destination specified by the subscription.
Last Message Time	The time when a message was last sent to the destination specified by the subscription.
Message Count	The number of messages put to the destination that are specified by this subscription.
Multicast Reliability Indicator	The multicast reliability indicator.
RateMessage Count	The rate of messages put to the destination specified by this subscription.
Resume Date	The date of the most recent MQSUB API call that connected to this subscription.
Resume Time	The time of the most recent MQSUB API call that connected to this subscription.
Subscription Type	The type of subscription, which indicates how the subscription was created.
Subscription User	The user ID that owns the subscription.
TopicString	The name of the topic string.
Expired	When checked, performance data has not been received in the time specified in the Duration region on the RTView Configuration Application > Solution Package Configuration > IBM MQ > DATA STORAGE tab.
Timestamp	The date and time of the last data update.

IBM MQ Subscriptions Heatmap





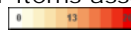

This display allows you to view current alert status and performance metrics based on the total subscriptions on a specific MQ broker in a heatmap format.

Each rectangle in the heatmap represents a different subscription, where the rectangle color indicates the most critical alert state for items associated with that subscription, and the rectangle size represents the tablespace allocation size for the subscription.

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.

Use the **Show Broker** 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 ["IBM MQ Subscription Summary"](#) display.

Alert Severity	<p>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:</p> <ul style="list-style-type: none">  (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	<p>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.</p>
Rate Message Count	<p>The rate of messages put to the destination specified by the subscription. The color 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 published topics in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>

IBM MQ Subscription Summary

This display allows you to view current alert status and performance metrics for a specific subscription on a specific MQ broker or all subscriptions on all brokers. Select a broker, a subscription, and a subscription ID from the drop-down menus to view data for the selected subscription. Clicking on the **Message Count** and **Rate Message Count** information boxes at the top of the display opens the “[IBM MQ Subscriptions Table](#)” display, where you can sort and compare the performance values of all subscriptions.

The trend graph traces the **Metric Trends (Rate Message Count)** for the selected subscription. 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. You can hover over the trend graph to see the values at a particular time.

IBM MQ Subscription Summary ▾
25-Oct-2019 11:15 1 Alert ✓ DATA

Broker: vmrh5-3_legacy ▾ Name: SYSTEM.DEFAULT.SUB ▾
 Subscription ID: 414d5120514d5f303030312020202020561cf54c06000010 ▾

Message Count
0

Rate Message Count
0.00

Metric Trends
Log Scale: 15 minutes ⌚

Critical/Warning: 1/0	Connected: true	Expired: false
Subscription ID: 414d5120514d5f303030312020202020561cf54c06000010		
ActiveConnection: 00		
DeltaMessage Count: 0	Deltatime_stamp: 1053	Durable: YES
Last Message Date:	Last Message Time:	Message Count: 0
Multicast Reliability Indicator: 0	RateMessage Count: 0.00	Resume Date:
Resume Time:	Subscription Type: ADMIN	Subscription User:
TopicString:		
Last Update: 25-Oct-2019 11:16:01		

Filter By:

- Broker:** Select the MQ broker for which you want to view data.
- Name** Select the name of the subscription for which you want to view data.
- Subscription ID** Select the subscription ID for which you want to view data.

Fields and Data

Message Count The number of messages put to the destination that are specified by this subscription.

Rate Message Count The rate of messages put to the destination that are specified by this subscription.

Metric Trends Graph

Traces the following for the selected topic:


Rate Message Count: Traces the rate of messages put to the destination specified by this subscription.

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

Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

Critical/Warning The total number of critical and warning alerts.

Connected When checked, denotes that the broker is connected.

Expired	When checked, performance data has not been received in the time specified in the Duration region on the RTView Configuration Application > Solution Package Configuration > IBM MQ > DATA STORAGE tab.
Subscription ID	The ID of the subscription.
Active Connection	The ConnId of the HConn that currently has this subscription open.
Delta Message Count	The number of messages put to the destination that are specified by this subscription since the last data update.
Last Message Date	The date that a message was last sent to the destination specified by the subscription.
Multicast Reliability Indicator	The multicast reliability indicator.
Resume Time	The time of the most recent MQSUB API call that connected to this subscription.
Topic String	The name of the topic string.
Deltatime_stamp	The difference in time (in milliseconds) between the most recent update (Last Update/ Timestamp) and the previous update.
Last Message Time	The time when a message was last sent to the destination specified by the subscription.
RateMessage Count	The rate of messages put to the destination specified by this subscription.
Subscription Type	The type of subscription, which indicates how the subscription was created.
Durable	When NO, the subscription is removed when the application that created it is closed or disconnected from the queue manager. When YES, the subscription persists even when the creating application is no longer running or has been disconnected. The subscription is reinstated when the queue manager is restarted.
Message Count	The number of messages put to the destination that are specified by this subscription.
Resume Date	The date of the most recent MQSUB API call that connected to this subscription.
Subscription User	The user ID that owns the subscription.
Last Update	The date and time of the last data update.

IBM MQ Subscription Configs View

The displays in this view allow you to view the current alert status and performance metrics for your subscription configurations on a specific MQ broker. Clicking **IBM MQ Subscription Configs** in the left/navigation menu opens the [“IBM MQ Subscription Configs Table”](#) display, which allows you to view current alert status and performance metrics of all subscription configurations on a specific MQ broker in a table format. The other available display in this View is:

- [“IBM MQ Subscription Config Summary”](#): Clicking **MQ Subscription Config Summary** in the left/navigation menu opens this display, which allows you to view current alert status and performance metrics of a subscription configuration on a specific MQ broker in a table format.

IBM MQ Subscription Configs Table

View current alert status and performance metrics of all subscription configurations on a specific MQ broker in a table format.

Each row in the table contains data for a particular subscription config. 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**.

You can also 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 subscription by double-clicking a row, which opens the ["IBM MQ Subscription Config Summary"](#) display.

IBM MQ Subscription Configs Table 25-Oct-2019 13:36 ✓ DATA

Broker: - All -

Subscription Objects: 44

All Subscription Configs Table

Broker	Name	Subscription ID	Alteration Date
vmrh5-2_legacy			
vmrh5-2_legacy			
vmrh5-2_legacy			
vmrh5-2_legacy			
vmrh5-2_legacy	SYSTEM.DEFAULT.SUB		
vmrh5-2_legacy	QM_0001 SYSTEM.BROKER.INTER.BROKER.COMMUNICATIONS 414D515901010000000000		
MQ8.0_TB33	SYSTEM.DEFAULT.SUB		
MQ8.0_TB33	UNSECURE8 SYSTEM.BROKER.INTER.BROKER.COMMUNICATIONS 414D5159010100000000		
vmrh5-1_legacy			
vmrh5-1_legacy			
vmrh5-1_legacy			

Page 1 of 2 1 - 40 of 44 items

Filter By:

Broker: Select the MQ broker for which you want to view subscription configs or select **All** to view all subscription configs for all brokers.

Count The number of topics found on the broker and listed in the table.

All Subscription Configs Table

Broker The name of the broker.

Name The name of the subscription config.

Subscription ID The internal, unique key identifying the subscription.

Alteration Date The date of the most recent MQSUB or Change Subscription command that modified the properties of the subscription.

Alteration Time	The time of the most recent MQSUB or Change Subscription command that modified the properties of the subscription.
Creation Date	The creation date of the subscription.
Creation Time	The creation time of the subscription.
Destination	Specifies the name of the alias, local, remote, or cluster queue to which messages for this subscription are put.
Destination Class	When MANAGED, the destination is managed. When PROVIDED, the destination queue is as specified in the Destination field.
Destination Correlation ID	A correlation identifier that is placed in the CorrelId field of the message descriptor for all the messages sent to this subscription.
Destination Queue Manager	The name of the destination queue manager, either local or remote, to which messages for the subscription are forwarded.
Durable	When NO, the subscription is removed when the application that created it is closed or disconnected from the queue manager. When YES, the subscription persists even when the creating application is no longer running or has been disconnected. The subscription is reinstated when the queue manager is restarted.
Expiry	The time, in tenths of a second, at which a subscription expires after its creation date and time. A value of unlimited means that the subscription never expires.
Publish Priority	The priority of messages sent to this subscription. Possible values are: AS PUBLISHED -- The priority of messages sent to this subscription is taken from that priority supplied to the published message. AS_QDEF -- The priority of messages sent to this subscription is determined by the default priority of the queue defined as a destination. 0-9 -- An integer value providing an explicit priority for messages sent to this subscription.
Publish Subscribe Properties	Specifies how publish/subscribe related message properties are added to messages sent to this subscription.
Published Accounting Tokens	The value of the accounting token used in the AccountingToken field of the message descriptor.
Published Application Identity Data	The value of the application identity data used in the ApplIdentityData field of the message descriptor.
Request Only	Indicates whether the subscriber polls for updates using the MQSUBRO API call, or whether all publications are delivered to this subscription. Possible values are: ALL -- All publications on the topic are delivered to this subscription. ON REQUEST -- Publications are only delivered to this subscription in response to an MQSUBRO API call.
Selector	Specifies the selector applied to messages published to the topic.
Selector Type	The type of selector string that has been specified.
Subscription Level	The level within the subscription interception hierarchy at which this subscription is made.

Subscription Scope	Lists whether this subscription is passed to other queue managers in the network. Possible values are: ALL -- The subscription is forwarded to all queue managers directly connected through a publish/subscribe collective or hierarchy. QUEUE MANAGER -- The subscription only forwards messages published on the topic within this queue manager.
Subscription Type	The type of subscription, which indicates how the subscription was created.
Subscription User	The user ID that owns the subscription.
Topic Object	The name of a previously defined topic object from which is obtained the topic name for the subscription.
Topic String	The name of the topic string.
Userdata	Denotes the user data associated with the subscription.
Variable User	Denotes whether a user other than the one who created the subscription (the user shown in SubscriptionUser) can take over the ownership of the subscription
Wildcard Schema	Lists the schema to be used when interpreting any wildcard characters contained in the TopicString.
Expired	When checked, performance data has not been received in the time specified in the Duration region on the RTView Configuration Application > Solution Package Configuration > IBM MQ > DATA STORAGE tab.
Timestamp	The date and time of the last data update.

IBM MQ Subscription Config Summary

View current alert status and performance metrics of a subscription configuration on a specific MQ broker in a table format. Select a broker, a subscription, and a subscription ID from the drop-down menus to view data for the selected subscription config. Clicking on the information boxes at the top of the display opens the ["IBM MQ Subscription Configs Table"](#) display, where you can sort and compare the performance values of all subscription configs.

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 Subscription Config Summary 25-Oct-2019 14:13 No Alerts DATA

Broker: MQ7.5_TB23 Name: SYSTEM.DEFAULT.SUB
 Subscription ID: 414d5120554e534543555245202020d933745706000010

Alteration Date 2016-06-29	Alteration Time 13:47:22	Creation Date 2016-06-29	Creation Time 13:47:22
--------------------------------------	------------------------------------	------------------------------------	----------------------------------

Critical/Warning: **0/0** Connected: **true** Expired: **false**

Destination: **-1** Destination Class: **PROVIDED**
 Destination Correlation ID: **00**
 Destination Queue Manager: Durable: **YES** Expiry: **-1**
 Publish Priority: **AS PUBLISHED** Publish Subscribe Properties: **MESSAGE**
 Published Accounting Token: **00**
 Published Application Identity Data: Request Only: **ALL** Selector:
 Selector Type: **NONE** Subscription Level: **1** Subscription Scope: **ALL**
 Subscription Type: **ADMIN** Subscription User: Topic Object:
 Topic String: Userdata: Variable User: **ANY**
 Wildcard Schema: **TOPIC**

Last Update: **25-Oct-2019 14:13:32**

Filter By:

- Broker:** Select the MQ broker for which you want to view data.
- Name** Select the name of the subscription configuration for which you want to view data.
- Subscription ID** Select the subscription ID for which you want to view data.

Fields and Data

- Alteration Date** The date of the most recent MQSUB or Change Subscription command that modified the properties of the subscription.
- Alteration Time** The time of the most recent MQSUB or Change Subscription command that modified the properties of the subscription.
- Creation Date** The creation date of the subscription.
- Creation Time** The creation time of the subscription
- Critical/Warning** The total number of critical and warning alerts.
- Connected** When checked, denotes that the broker is connected.
- Expired** When checked, performance data has not been received in the time specified in the **Duration** region on the RTView Configuration Application > **Solution Package Configuration** > **IBM MQ** > **DATA STORAGE** tab.

Destination	Specifies the name of the alias, local, remote, or cluster queue to which messages for this subscription are put.
Destination Correlation ID	A correlation identifier that is placed in the CorrelId field of the message descriptor for all the messages sent to this subscription.
Destination Queue Manager	The name of the destination queue manager, either local or remote, to which messages for the subscription are forwarded.
Publish Priority	The priority of messages sent to this subscription. Possible values are: AS PUBLISHED -- The priority of messages sent to this subscription is taken from that priority supplied to the published message. AS_QDEF -- The priority of messages sent to this subscription is determined by the default priority of the queue defined as a destination. 0-9 -- An integer value providing an explicit priority for messages sent to this subscription.
Published Accounting Token	The value of the accounting token used in the AccountingToken field of the message descriptor.
Published Application Identity Data	The value of the application identity data used in the ApplIdentityData field of the message descriptor.
Selector Type	The type of selector string that has been specified.
Subscription Type	The type of subscription, which indicates how the subscription was created.
Topic String	The name of the topic string.
Wildcard Schema	Lists the schema to be used when interpreting any wildcard characters contained in the TopicString.
Destination Class	When MANAGED , the destination is managed. When PROVIDED , the destination queue is as specified in the Destination field.
Durable	When NO , the subscription is removed when the application that created it is closed or disconnected from the queue manager. When YES , the subscription persists even when the creating application is no longer running or has been disconnected. The subscription is reinstated when the queue manager is restarted.
Publish Subscribe Properties	Specifies how publish/subscribe related message properties are added to messages sent to this subscription.
Request Only	Indicates whether the subscriber polls for updates using the MQSUBRQ API call, or whether all publications are delivered to this subscription. Possible values are: ALL -- All publications on the topic are delivered to this subscription. ON REQUEST -- Publications are only delivered to this subscription in response to an MQSUBRQ API call.
Subscription Level	The level within the subscription interception hierarchy at which this subscription is made.
Subscription User	The user ID that owns the subscription.
Userdata	Denotes the user data associated with the subscription.
Selector	Specifies the selector applied to messages published to the topic.
Subscription Scope	Lists whether this subscription is passed to other queue managers in the network. Possible values are: ALL -- The subscription is forwarded to all queue managers directly connected through a publish/subscribe collective or hierarchy. QUEUE MANAGER -- The subscription only forwards messages published on the topic within this queue manager.

Topic Object	The name of a previously defined topic object from which is obtained the topic name for the subscription.
Variable User	Denotes whether a user other than the one who created the subscription (the user shown in SubscriptionUser) can take over the ownership of the subscription
Last Update	The date and time of the last data update.

IBM WebSphere

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"](#)
- ["WebSphere Servers Heatmap"](#): Performance metrics for one IBM WebSphere Server, including current and historic performance metrics.
- ["WebSphere Server Summary"](#): Heatmap of performance metrics for all Web modules for one IBM WebSphere Server.
- ["WebSphere Apps Table"](#): Table and trend graphs of performance metrics for Web modules.
- ["WebSphere Apps Heatmap"](#)
- ["WebSphere App Summary"](#)

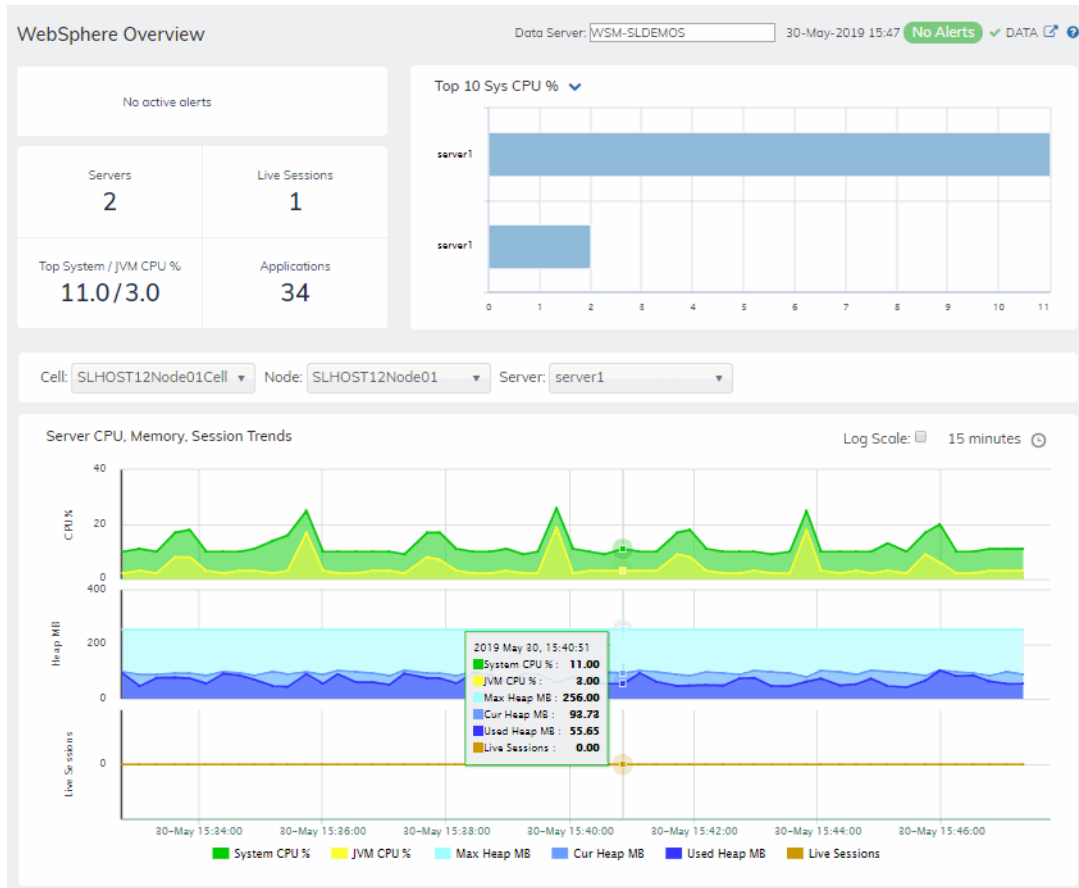
WebSphere Overview

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

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"** display and view details for that server.

WebSphere Servers Heatmap

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.

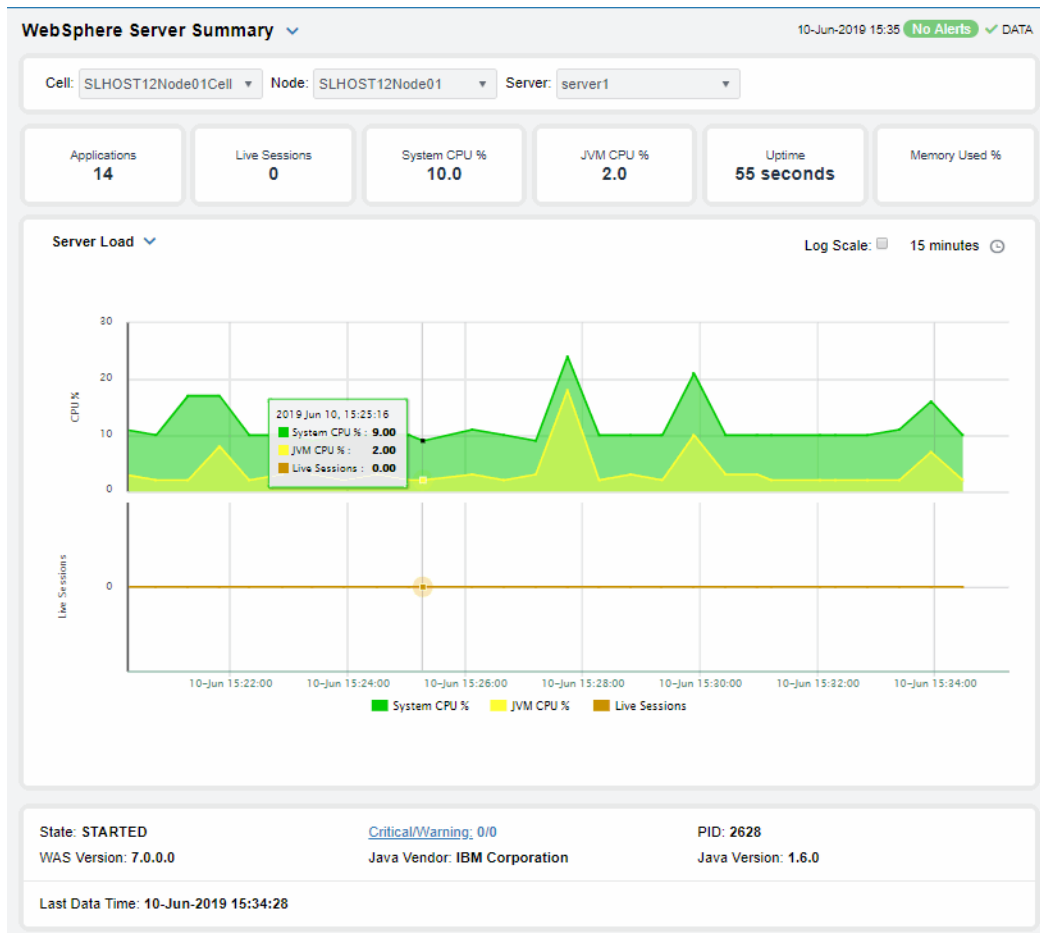
Use the available drop-down menus or right-click to filter data shown in the display. Use the check-boxes to include or exclude labels in the heatmap. Move your mouse over a rectangle to see additional information. Drill-down and investigate by clicking a rectangle in the heatmap to view details for the selected Web module in the **Server Summary** display.

WebSphere Server Summary

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

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

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

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.

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”](#)
- [“Docker”](#)
- [“JBoss”](#)
- [“MongoDB”](#)
- [“MySQL Database”](#)
- [“MS SQL”](#)
- [“Node.js”](#)
- [“RTView Host Agent”](#)
- [“VMware vCenter”](#)

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 following Solution Package for Amazon Web Services features an overview display, [“Amazon EC2 Overview”](#) (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”](#)
- [“Amazon EC2 Instance Table”](#)
- [“Amazon EC2 Instance Summary”](#)

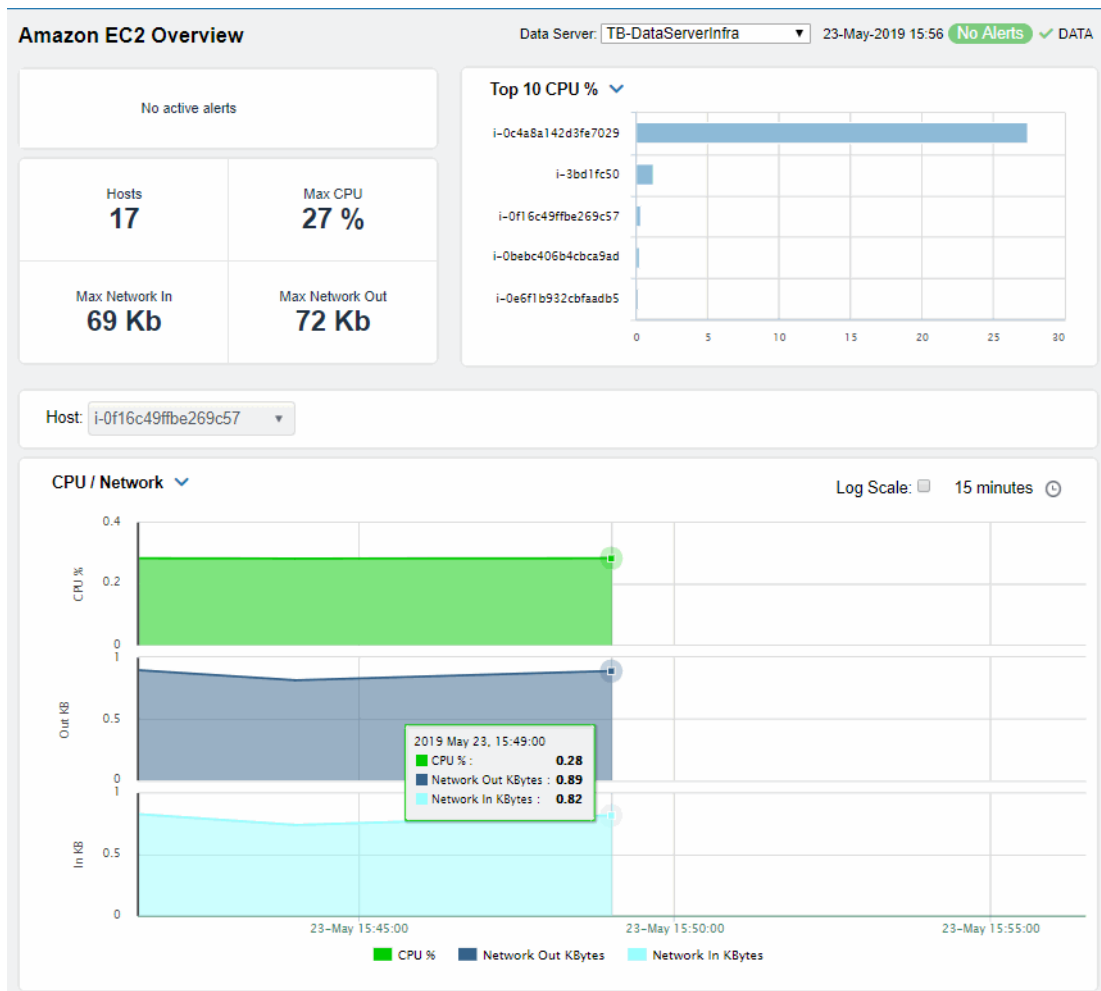
Amazon EC2 Overview

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.



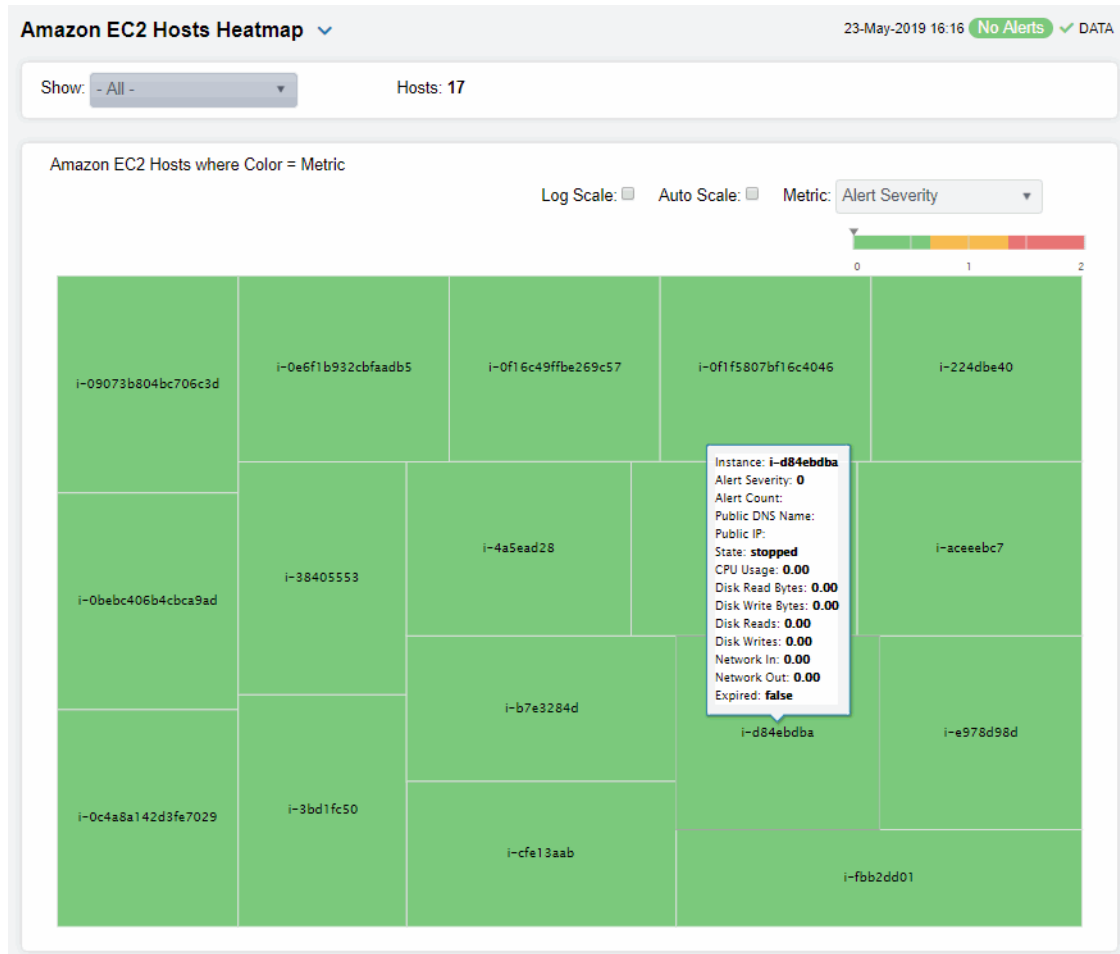
Amazon EC2 Instance Heatmap

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”](#) 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.

Alert Severity	<p>The maximum level of alerts in the rectangle. Values range from 0 - 2, as indicated in the color gradient  bar, where 2 is the highest Alert Severity:</p> <ul style="list-style-type: none">  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.
CPU Usage	<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 0 to the maximum amount in the heatmap. The middle value in the gradient bar indicates the average amount.</p>
Network In	<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 0 to the maximum amount in the heatmap. The middle value in the gradient bar indicates the average amount.</p>
Network Out	<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 0 to the maximum amount in the heatmap. The middle value in the gradient bar indicates the average amount.</p>
Disk Reads	<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 0 to the maximum amount in the heatmap. The middle value in the gradient bar indicates the average amount.</p>
Disk Writes	<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 0 to the maximum amount in the heatmap. The middle value in the gradient bar indicates the average amount.</p>
Disk Read Bytes	<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 0 to the maximum amount in the heatmap. The middle value in the gradient bar indicates the average amount.</p>
Disk Write Bytes	<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 0 to the maximum amount in the heatmap. The middle value in the gradient bar indicates the average amount.</p>

Amazon EC2 Instance Table

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.

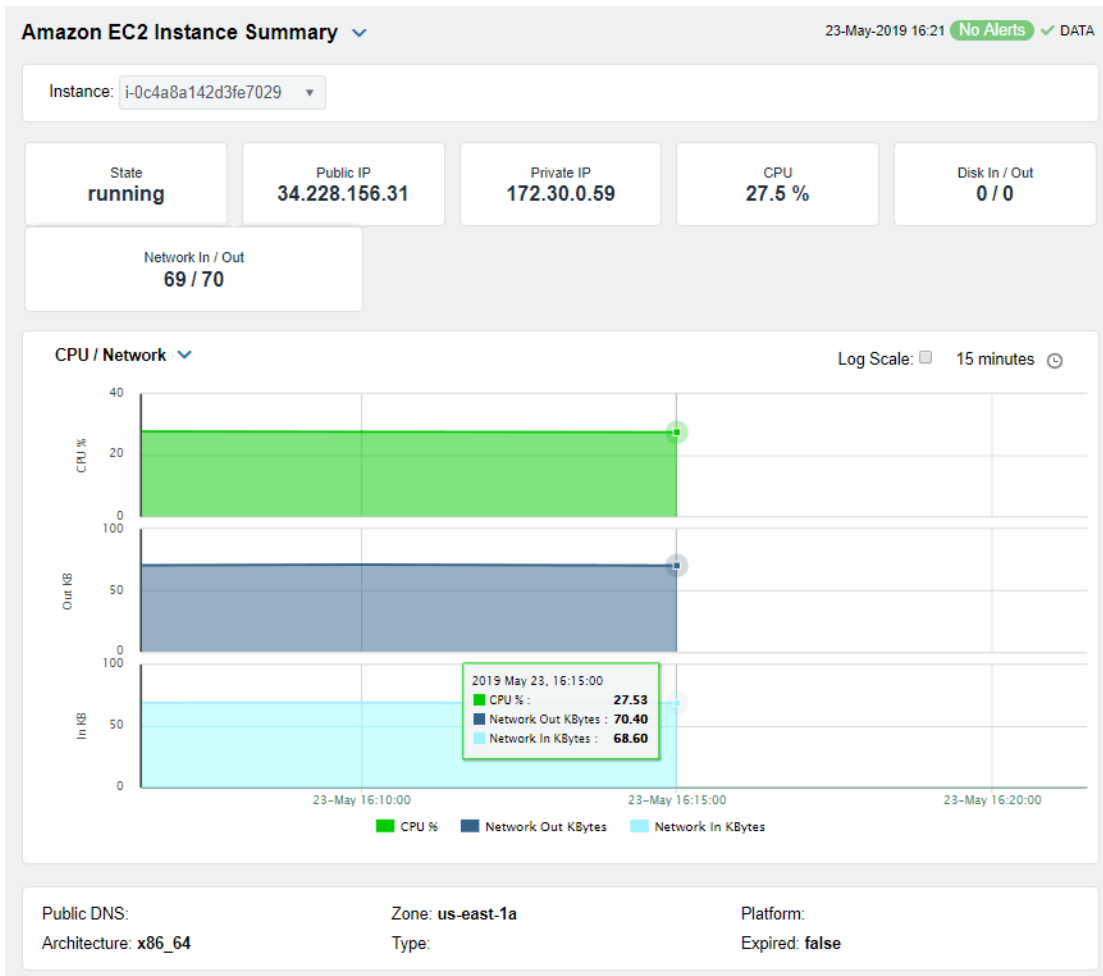
State	The instance state (running or stopped).
%CPU	The percent CPU used.
Disk Reads (bytes)	The amount of disk reads, in bytes.
Disk Reads (Ops)	The number of disk reads (count).
Disk Writes (bytes)	The amount of disk writes, in bytes.
Disk Writes (Ops)	The number of disk writes (count).
Network In	The number of incoming bytes.
Network Out	The number of outgoing bytes.
Instance Type	The instance type (e.g. m1.small).
Private IP	The instance private IP address.
Public IP	The instance public IP address.
Platform	The instance operating system (e.g. windows).
Architecture	The instance architecture (e.g. i386).
Image ID	The unique identifier for the image. For details about Amazon EC2 data, refer to vendor documentation.
Root Device Name	The name of the root device. For details about Amazon EC2 data, refer to vendor documentation.
Root Device Type	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

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 following Docker Views and their associated displays are available in the Monitor. This section describes the Monitor displays and includes:

- [“Docker Overview”](#): Describes the **Docker Overview** display.
- [“Docker Engines 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.
- [“Docker Containers 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.

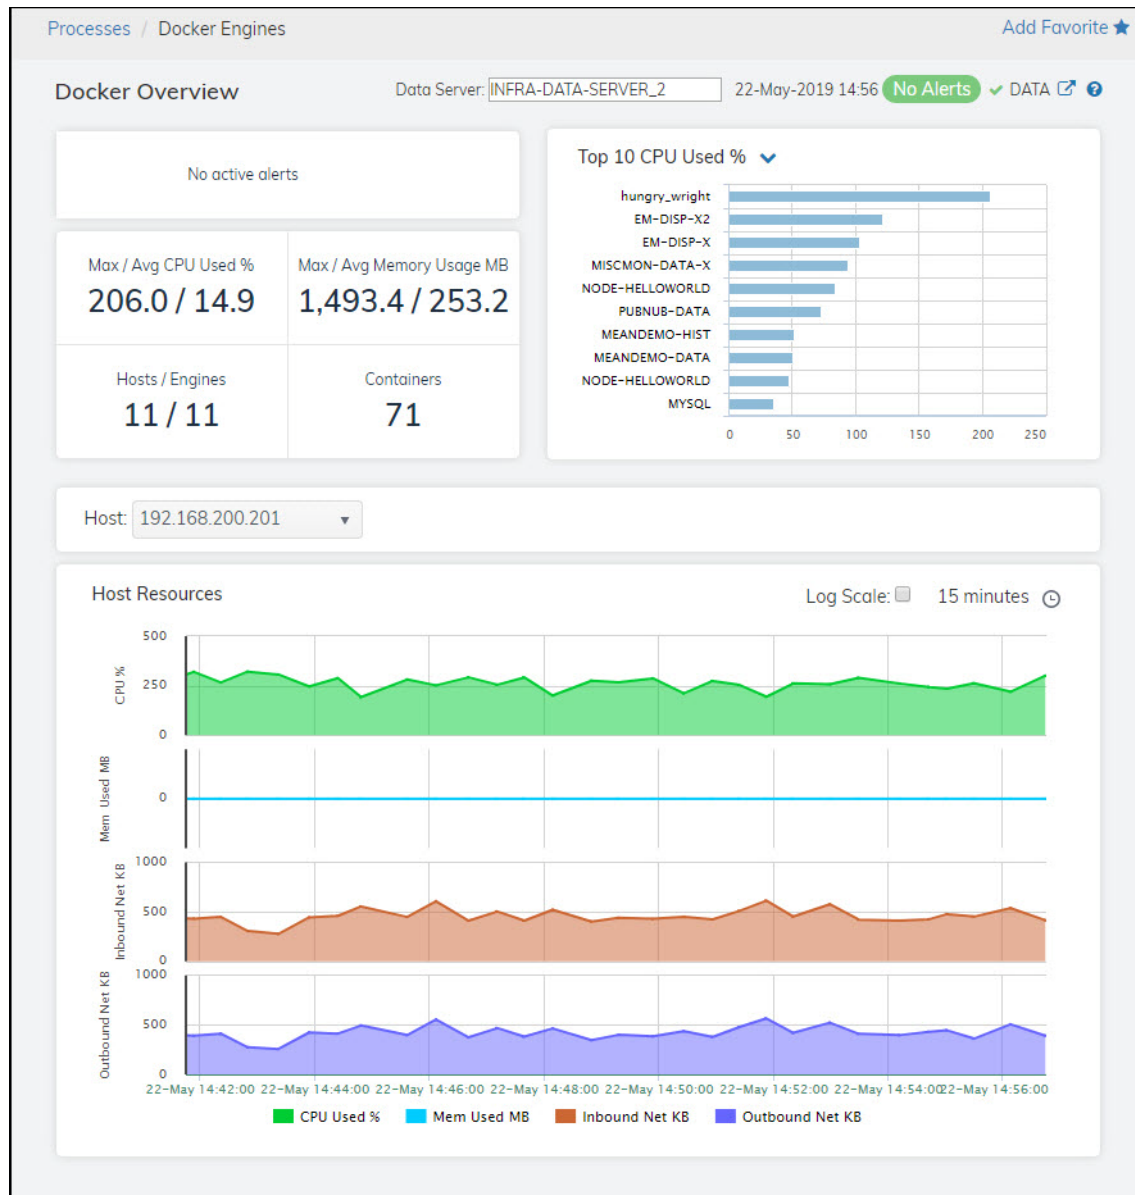
Docker Overview

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

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](#)” 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](#)” display, which provides a heatmap view of all engines and their associated metrics.
- **Single Engine Summary:** Opens the “[Docker Single Engine Summary](#)” display, which provides additional details and a way to view trending data for a single engine.

Docker 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. 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](#)” 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

Fields and Data:

Engines The total number of engines being monitored and listed in the table.

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.

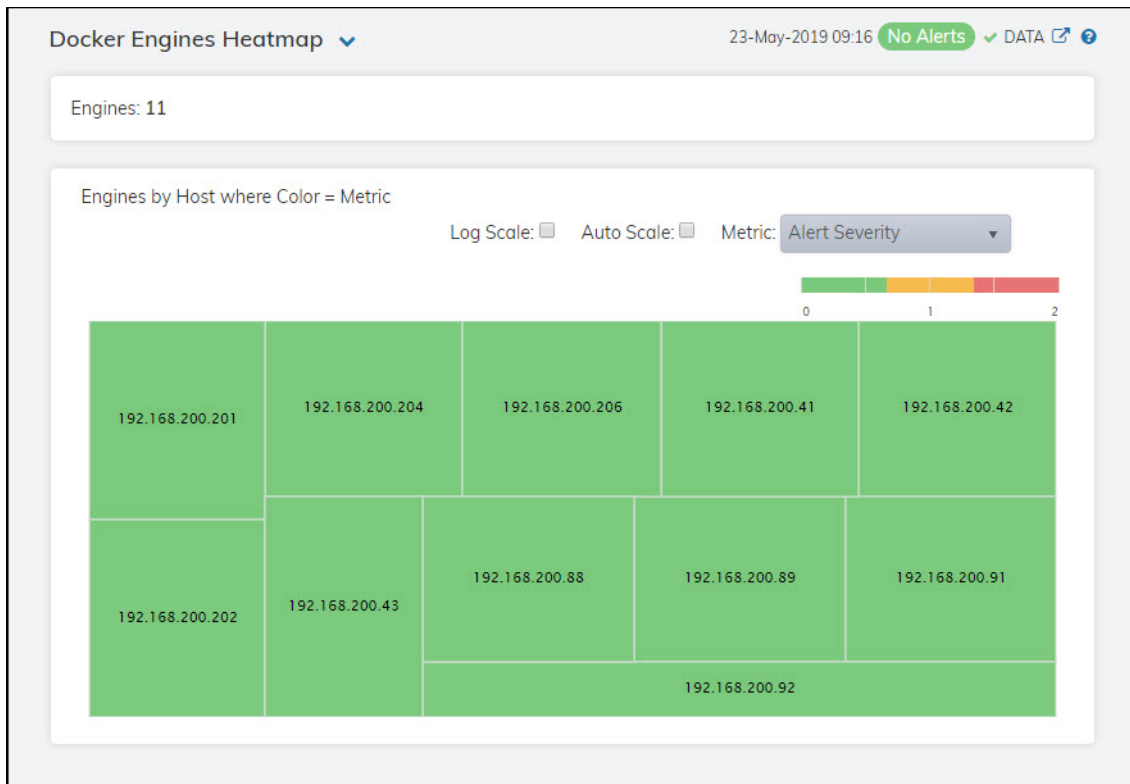
CPU Usage %	The percentage of CPU used by the engine.
Memory Avg MB	The average amount of memory, in megabytes, that is available to the engine.
Memory Use MB	The current memory usage by the engine, in megabytes, which includes all memory regardless of when it was accessed.
Memory WS MB	The amount of memory (in megabytes) in the working set, which includes recently accessed memory, dirty memory, and kernel memory.
Memory RSS MB	The amount of anonymous and swap cache memory (including transparent/hugepages), in megabytes.
Memory Fails	The number of times when enough memory has not been available.
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.
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.
Time Stamp	The date and time the row data was last updated.

Docker Engines Heatmap

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"](#) display.



Fields and Data:

Engines Lists the total number of engines found using the search parameters.










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.

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 Scale** is not selected.

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 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"](#) display for a detailed view of metrics for that particular engine.

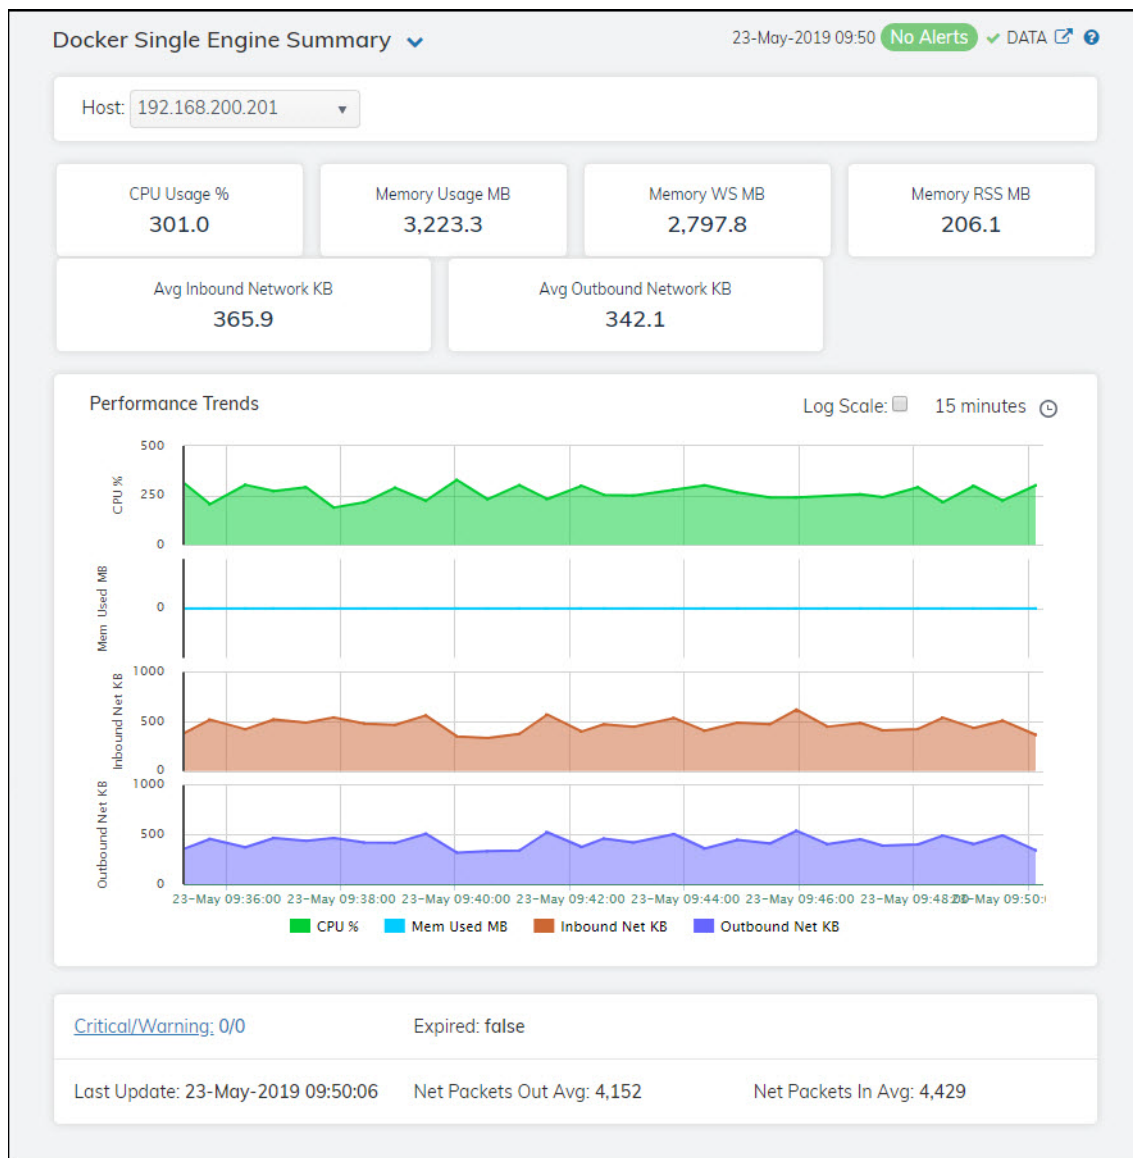
Alert Severity	<p>The current alert severity. Values range from 0 - 2, as indicated in the color gradient  bar, where 2 is the highest Alert Severity:</p> <ul style="list-style-type: none">  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	<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 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.</p>
CPU Usage	<p>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.</p> <p>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.</p>
Memory Usage MB	<p>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.</p> <p>The Auto Scale option does not impact this metric.</p>
Net Inbound Avg KB	<p>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.</p> <p>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.</p>
Net Outbound Avg KB	<p>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.</p> <p>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.</p>

Docker Single Engine Summary

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”](#) 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.



Filter By:

Host Select the host for which you want to show data in the display.

Fields and Data:

CPU Usage % The percentage of CPU used by the engine.

Memory Usage MB The current memory usage by the engine, in megabytes, which includes all memory regardless of when it was accessed.

Memory WS MB The amount of memory (in megabytes) in the working set, which includes recently accessed memory, dirty memory, and kernel memory.

Memory RSS MB The Resident Set Size, which is the amount of anonymous and swap cache memory (including transparent/hugepages), in kilobytes.

Avg Inbound Network KB The average number of incoming kilobytes.

Avg Outbound Network KB The average number of outbound kilobytes.

Performance Trends Graph

Traces the following:

CPU % -- traces the percentage of CPU being used on the engine.


Mem Used MB -- traces the amount of memory, in megabytes, used by the engine.

Inbound Net KB -- traces the average number of incoming kilobytes per second.

Outbound Net KB -- traces the average number of outgoing 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.

Time Settings Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

- Critical/Warning** The number of critical and warning alerts.

- 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.

- Last Update** The date and time of the last data update.

- Net Packets In Avg** The average number of incoming packets.

- Net Packets Out Avg** The average number of outgoing packets.

Docker Containers 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. Clicking **Docker Containers** from the left/navigation menu opens the “[Docker Containers Table](#)” 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](#)” 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](#)” display, which provides additional details and a way to view trending data for a single container for a particular host.

Docker Containers Table

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](#)” display by double-clicking on a row in the resulting table.

Docker Containers Table 23-May-2019 11:10 No Alerts DATA

Host: 192.168.200.201

Containers: 11

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	53c89d6f7f350e	✓		93.66	
192.168.200.201	MONGO-1	439618f9f2c16e	✓		0.47	
192.168.200.201	MONGO-RS-1	fd28c47fb254e6	✓		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	9fbad85f9f4286	✓		0.79	

Filter By:

The display includes these filtering options:

Host Select the host for which you want to show data in the display.

Containers Lists the total number of containers (rows) found using the search parameters.




All Containers Table

Host The name of the host.

Container The name of the 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 Avg MB The average amount of memory, in megabytes, that is available to the container.

Memory Usage MB Current memory usage by the container, in megabytes, which includes all memory regardless of when it was accessed.

Memory WS MB The amount of memory (in megabytes) in the working set, which includes recently accessed memory, dirty memory, and kernel memory.

Memory RSS MB The Resident Set Size, which is the amount of anonymous and swap cache memory (including transparent/hugepages), in megabytes.

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.

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.

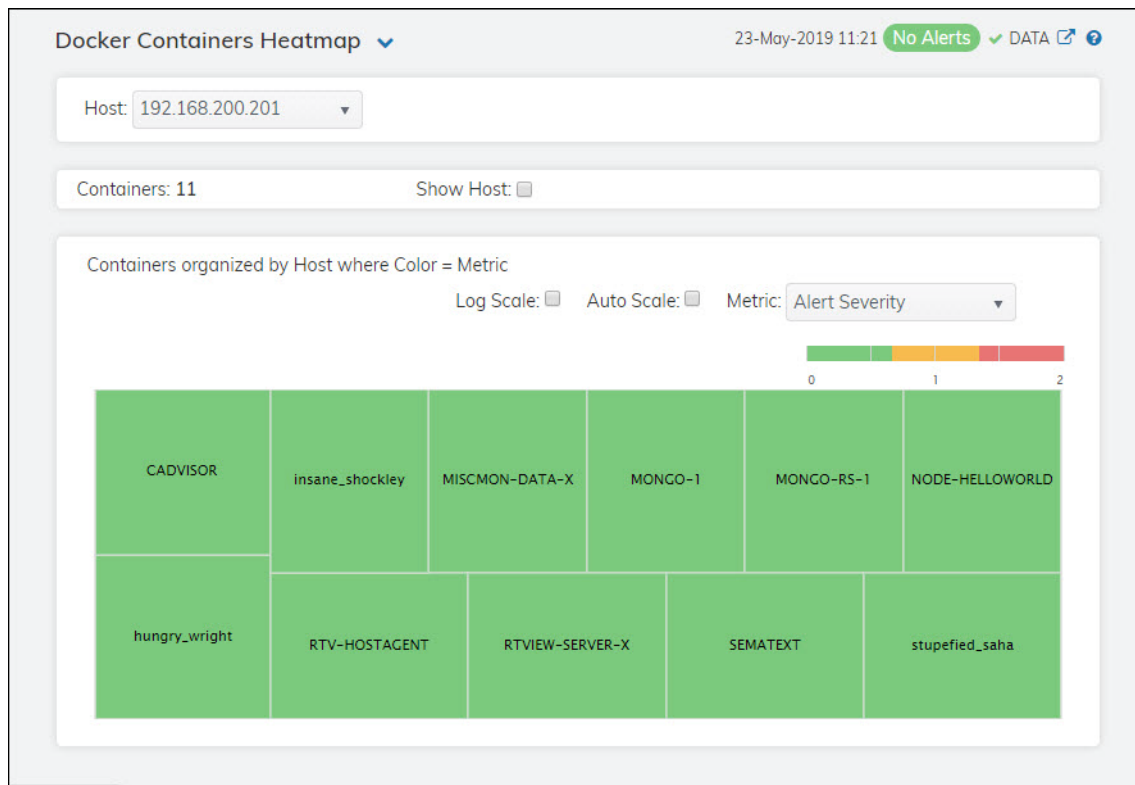
Timestamp The date and time the row data was last updated.

Docker Containers Heatmap

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”](#) display.



Filter:

Host Select the host (or **All Hosts**) for which you want to show data in the heatmap.

Fields and Data:

Count Lists the total number of containers (rows) found using the search parameters.

Show Host Select this check box to display the names of the containers at the top of each rectangle in the heatmap.

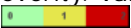
Heatmap




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.


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 Scale** is not selected.

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 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"](#) 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 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 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 Scale** 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 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.

**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 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.

Docker Single Container Summary

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”](#) 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.



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.

Fields and Data:

- CPU Usage %** The percentage of CPU used by the container.

Memory Usage MB	The current memory usage by the container, in megabytes, which includes all memory regardless of when it was accessed.
Memory WS MB	The amount of memory (in megabytes) in the working set, which includes recently accessed memory, dirty memory, and kernel memory.
Memory RSS MB	The Resident Set Size, which is the amount of anonymous and swap cache memory (including transparent/hugepages), in kilobytes.
Avg Inbound Network KB	The average number of incoming network kilobytes.
Avg Outbound Network KB	The average number of outbound network kilobytes.
Performance Trends Graph	<p>Traces the following:</p> <ul style="list-style-type: none"> CPU %-- traces percentage of CPU used by the container. Mem Used MB -- traces the current memory usage by the container, in megabytes, which includes all memory regardless of when it was accessed. Inbound Net KB -- traces the average number of incoming kilobytes per second. Outbound Net KB -- traces the average number of outgoing kilobytes per second. <p>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.</p>

Time Settings Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

ID	The name of the container.
Critical/Warning	The number of critical and warning alerts.
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.
Container ID	The absolute container name.
Last Update	The date and time of the last data update.
Net Packets Out Avg	The average number of outbound net packets.
Net Packets In Avg	The average number of inbound net packets.

JBoss

The following Solution Package for Red Hat JBoss features an overview display, "[JBoss Overview](#)" (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](#)"
- "[JBoss Applications](#)"

JBoss Overview

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

Displays in this View are:

- "[JBoss Servers Table](#)"
- "[JBoss Servers Heatmap](#)"
- "[JBoss Single Server Summary](#)"

JBoss Servers Table

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”](#) display and view details for that server.

JBoss Servers Heatmap

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”](#) display.

JBoss Single Server Summary

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”](#) 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

Displays in this View are:

- [“JBoss Apps Table”](#)
- [“JBoss Apps Heatmap”](#)
- [“JBoss App Totals”](#)

JBoss Apps Table

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”](#) display.

JBoss Apps Heatmap

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”](#) 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

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”](#) 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

This section describes the Solution Package for MongoDB which features an overview display, “[MongoDB Overview](#)” (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”](#)
- [“Mongo Databases View”](#)
- [“Mongo Collections View”](#)

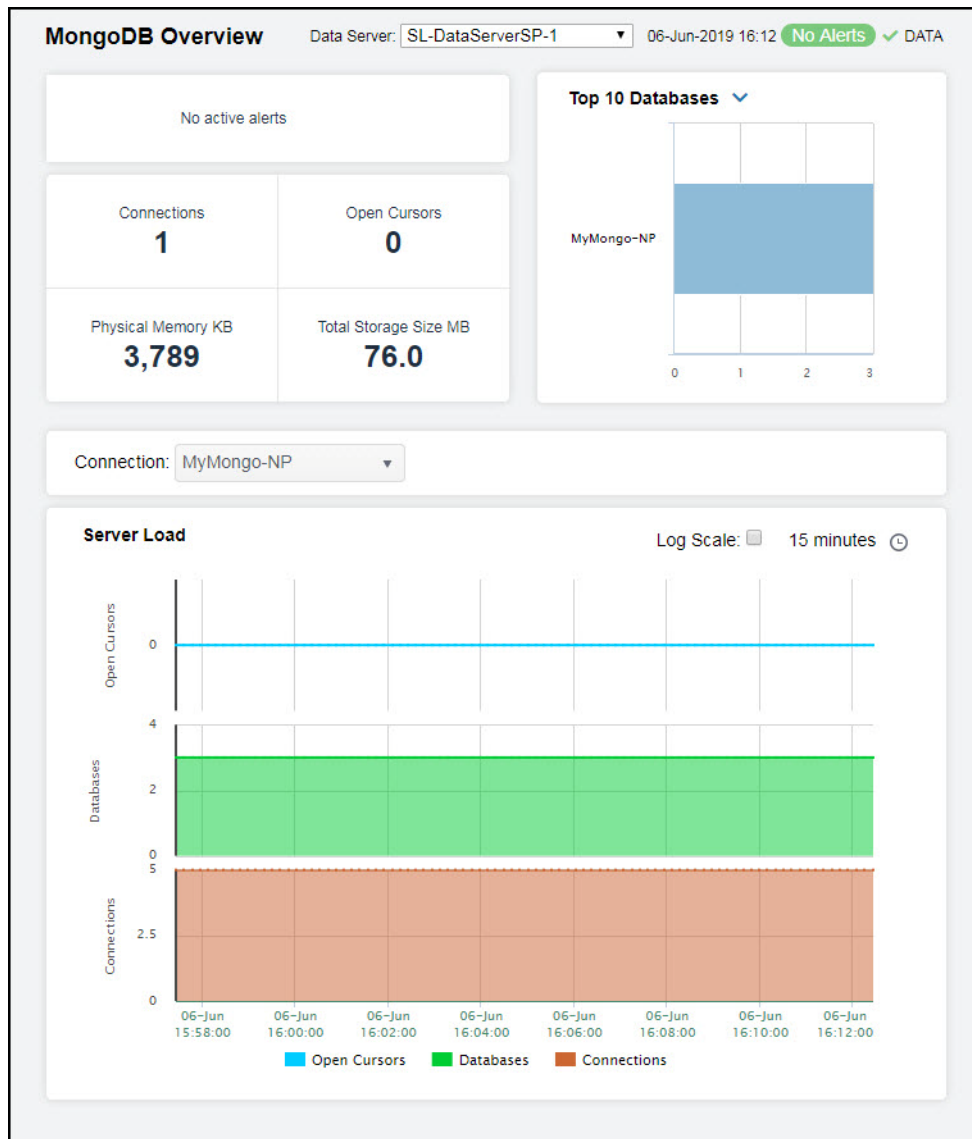
MongoDB Overview

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

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](#)” display, which shows all available utilization metrics for all MongoDB instances. The following displays are available:

- **All Instances Heatmap:** Opens the “[MongoDB Instances Heatmap](#)” display, which shows status and alerts for all MongoDB instances in a heatmap format.
- **Single Instance Summary:** Opens the “[MongoDB Instance Summary](#)” display, which enables you to view available utilization metrics for a single MongoDB instance.

MongoDB Instances Table

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”](#) 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	

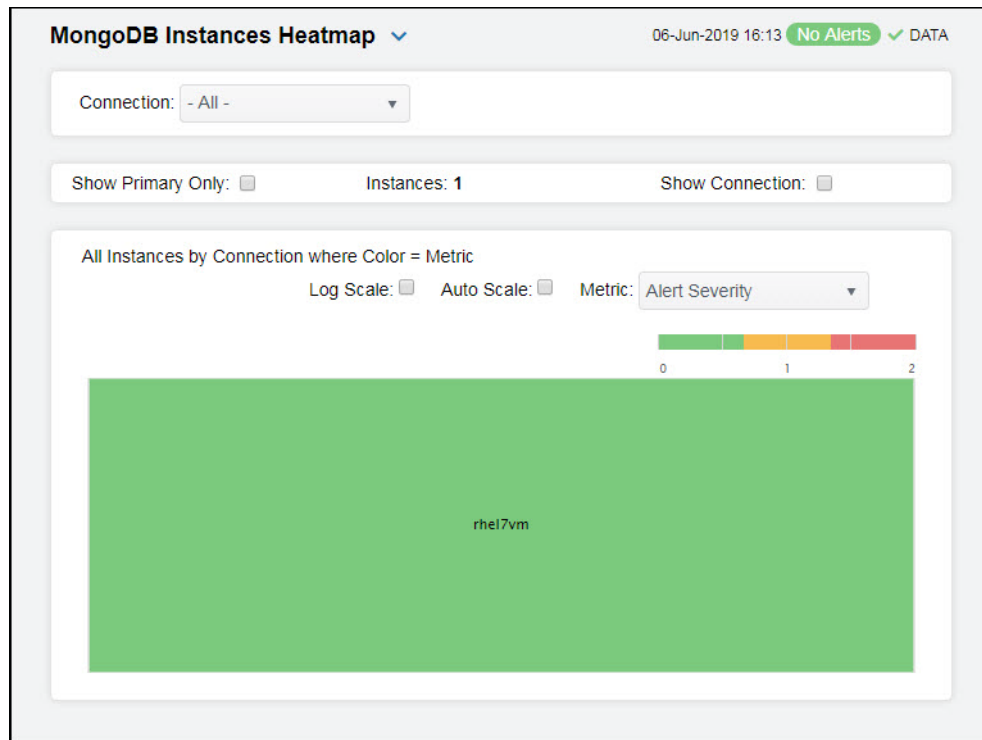
MongoDB Instances Heatmap

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”](#) 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.

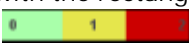
You can toggle between the commonly accessed displays by clicking the drop down list on the display title. Mouse-over rectangles to view more details about host performance and status.



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"](#) 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.



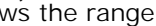
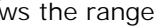
1 -- Metrics that have exceeded their specified **WARNINGLEVEL** threshold and have an Alert Severity value of **1** are shown in yellow. For a given rectangle, this indicates that one or more metrics have exceeded their warning threshold.

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	<p>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.</p> <p>The Auto option does not impact this metric.</p>
Open Cursors	<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 0 to the alert threshold of Mongol instanceOpenCursorsHigh, which is 2000. The middle value in the gradient bar indicates the middle value of the range (the default is 1000).</p> <p>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.</p>
Connections	<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 0 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 Auto option does not impact this metric.</p>
Databases	<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 0 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 Auto option does not impact this metric.</p>

MongoDB Instance Summary

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"](#) 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

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”](#) 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”](#) display, which shows status and alerts for all MongoDB databases in a heatmap format.
- **Single Database Summary:** Opens the [“MongoDB Single Database Summary”](#) 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

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”](#) 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

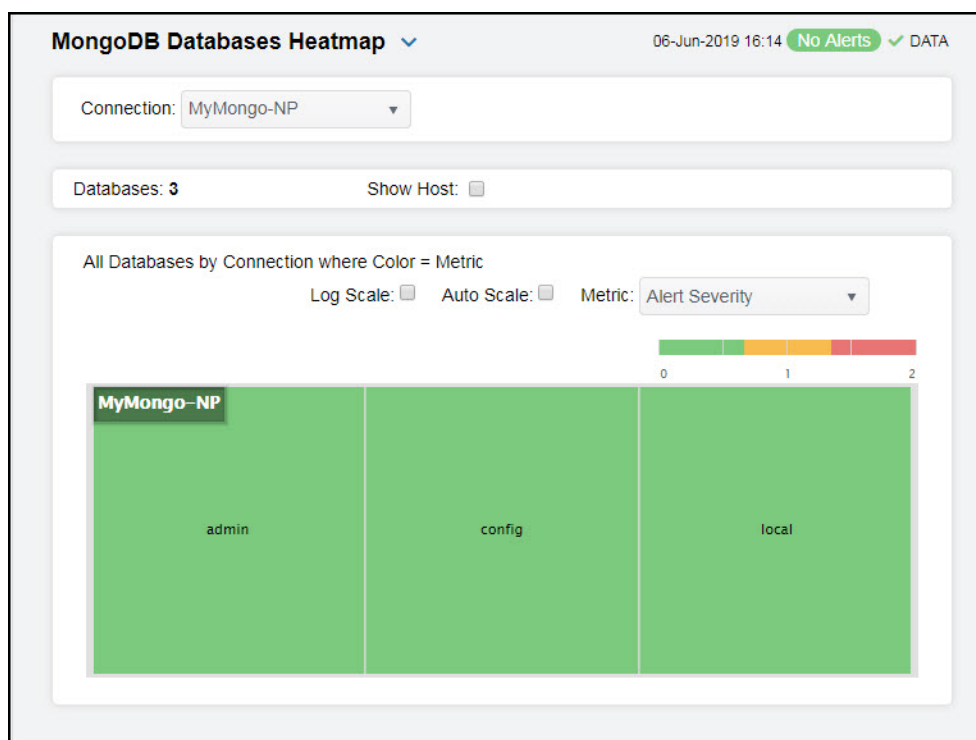
Connection	Host	Name	Alert Level	Alert Count	Status Up Down
MyMongo-NP	rhel7vm	admin	✓		✓
MyMongo-NP	rhel7vm	config	✓		✓
MyMongo-NP	rhel7vm	local	✓		✓

MongoDB Databases Heatmap

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"](#) 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"](#) display for a detailed view of metrics for that particular database.


Alert Severity	<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 2 is the greatest Alert Severity.</p> <p>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.</p> <p>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.</p> <p>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.</p>
Alert Count	<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 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>
Collections	<p>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.</p> <p>The Auto Scale option does not impact this metric.</p>
Data Size	<p>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).</p> <p>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.</p>

MongoDB Single Database Summary

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"](#) 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

These displays present several views of performance metrics for collections. Clicking **Mongo Collections** from the left/navigation menu opens the [“MongoDB Collections Table”](#) 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”](#) 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”](#) display, which shows detailed performance and utilization metrics and trends for a specified collection on a particular database.

MongoDB Collections Table

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”](#) display and view metrics for that particular collection. Toggle between the commonly accessed displays by clicking the drop down list on the display title.

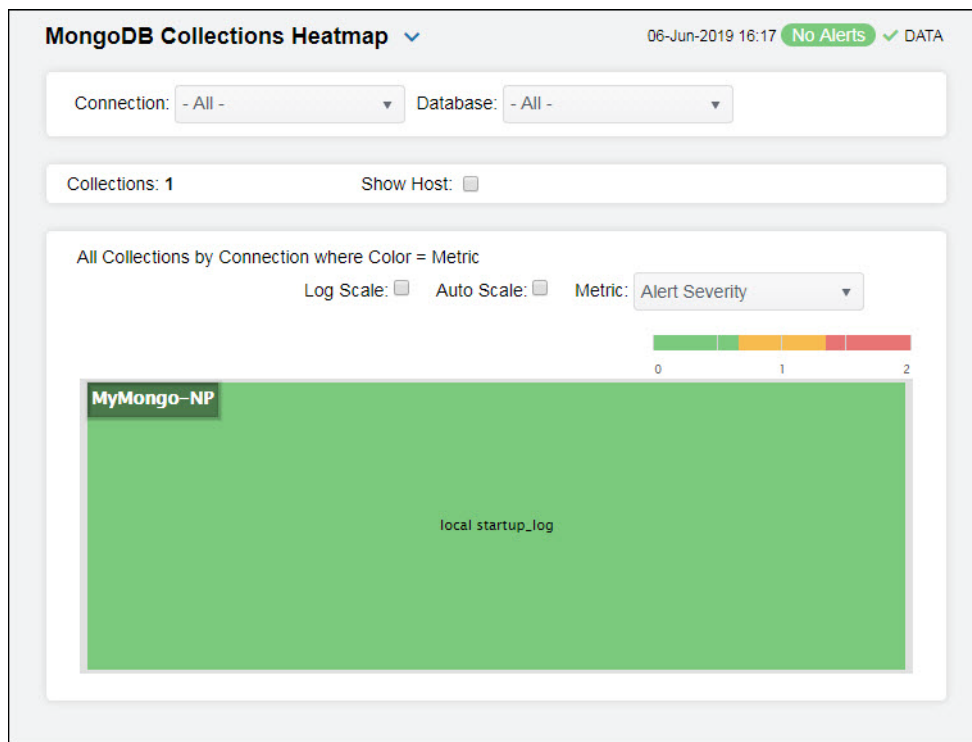
Connection	Host	Database	Name	Alert Level	Alert Count
MyMongo-NP	rhel7vm	local	startup_log	✓	

MongoDB Collections Heatmap

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"](#) 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"](#) display for a detailed view of metrics for that particular collection.

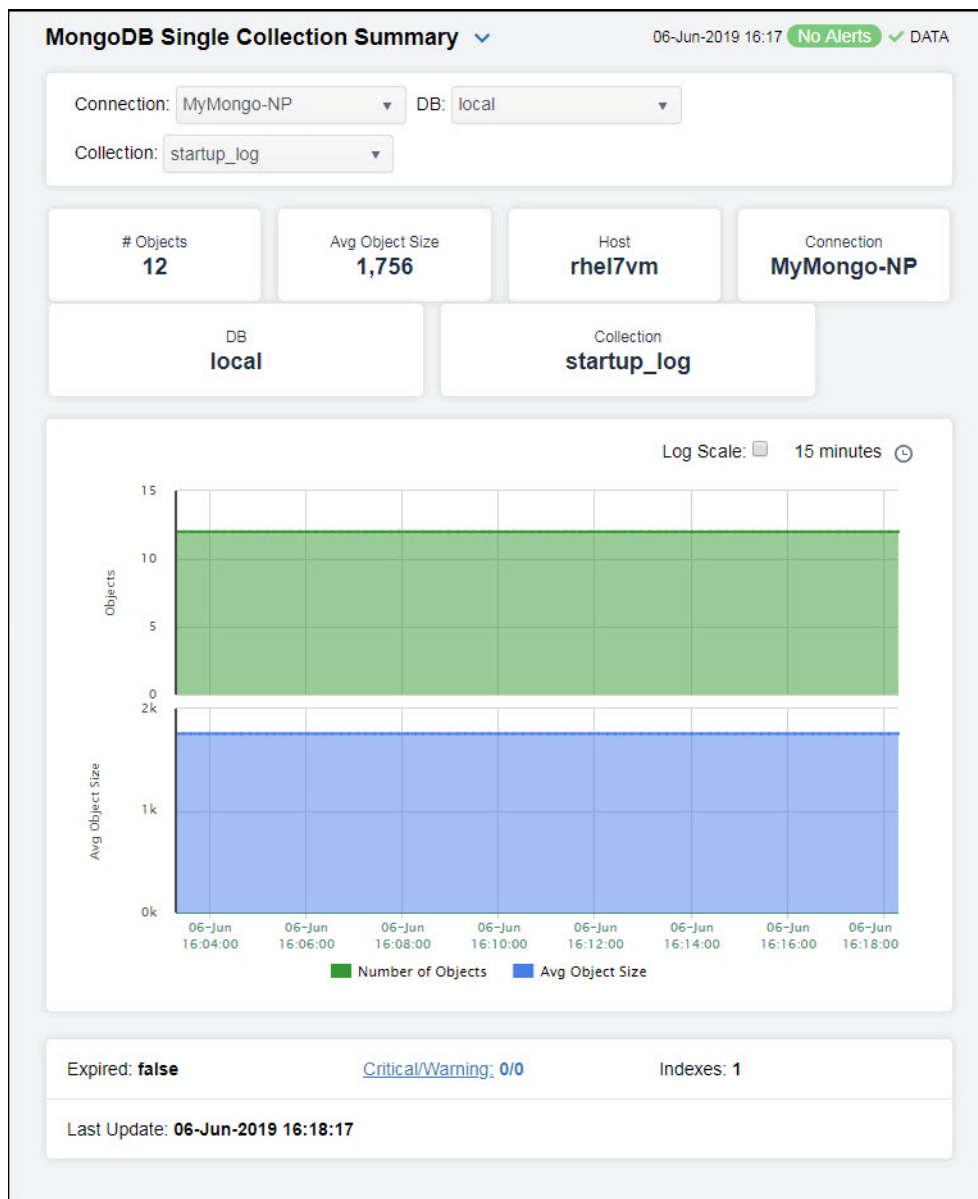
Alert Severity	<p>The maximum alert level in the item (index) associated with the rectangle. Values range from 0 to 2, as indicated in the color gradient bar , where 2 is the greatest Alert Severity.</p> <p>2 -- Metrics that have exceeded their specified ALARMLEVEL threshold and have an Alert Severity value of 2 are shown in red. For a given rectangle, this indicates that one or more metrics have exceeded their alarm threshold.</p> <p>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.</p> <p>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.</p>
Alert Count	<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 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>
Objects	<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 0 to the alert threshold of MongoCollectionNumObjectsHigh. The middle value in the gradient bar indicates the middle value of the range.</p> <p>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.</p>
Avg Object Size	<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 0 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

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"** 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 HTML displays provide extensive visibility into the health and performance of the MySQL database. The HTML version features an overview display, "[MySQL Overview Display](#)" (pictured below), and the following Views which can be found under **Components** tab > **Databases** > **MySQL Database**:

- "[MySQL Overview Display](#)"
- "[All MySQL Instances View](#)"
- "[Single MySQL Instance View](#)"

MySQL Overview Display

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.

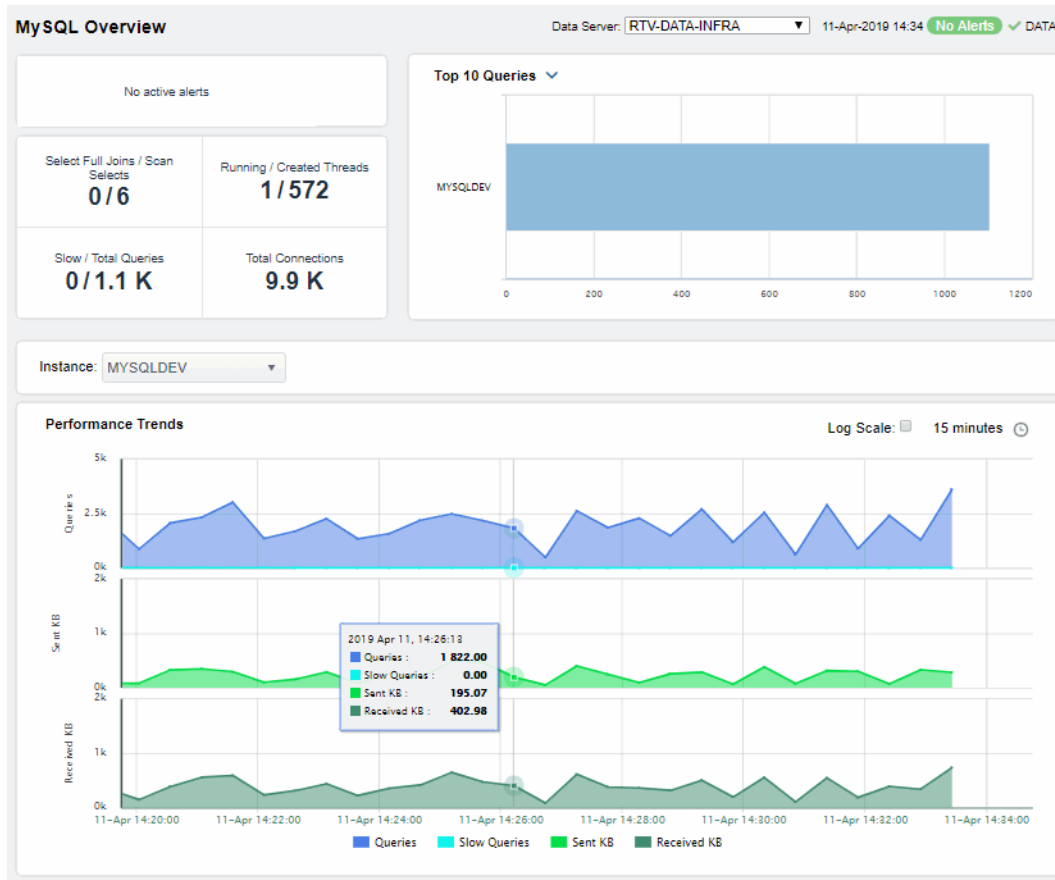
The total pending messages, the outgoing messages per second, and the incoming messages per second for a selected EMS Server on your connected DataServer.

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

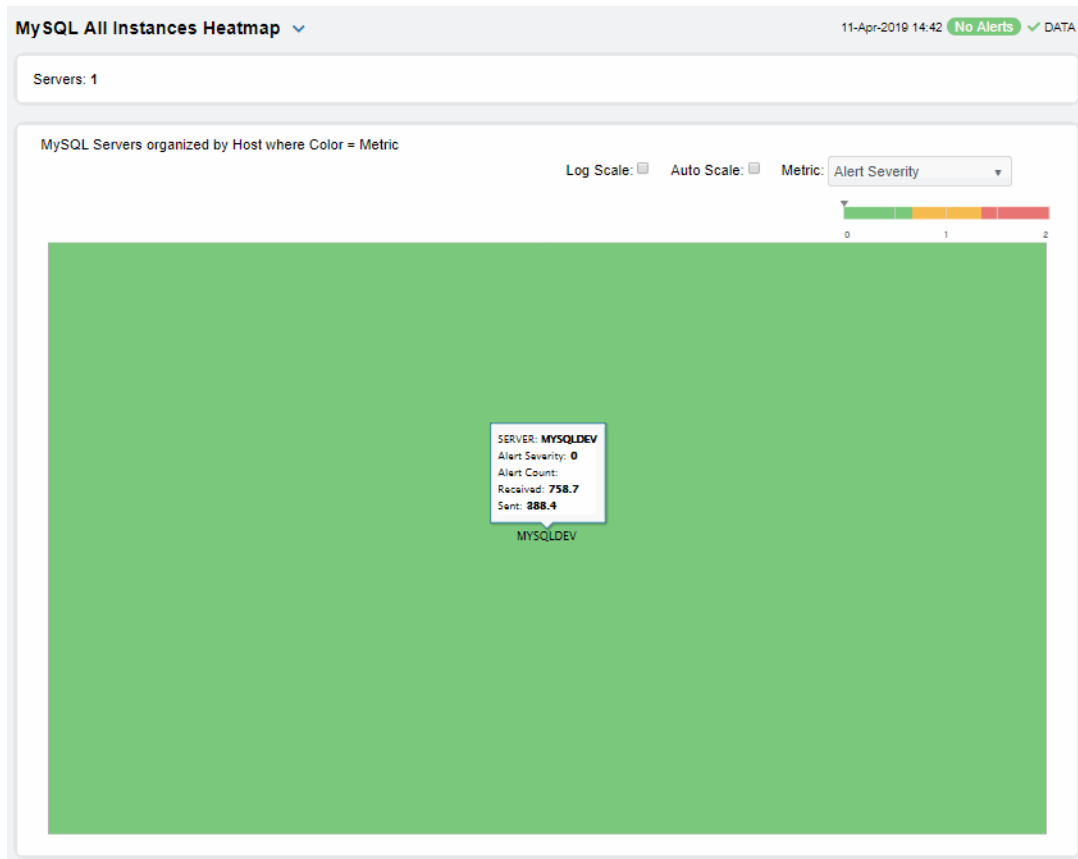
Displays in this View are:

- [“Instances Heatmap”](#): A heatmap view of all servers and their associated metrics.
- [“All MySQL Instances”](#): A tabular view of your servers and their associated metrics.

Instances Heatmap





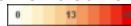


This heatmap display provides an easy-to-view interface that allows you to quickly identify the current status of each of your MySQL instances. Choose a metric from the **Metric** drop down menu. By default, this display shows the heatmap based on the **Alert Severity** metric. Other metrics are **Alert 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”](#) 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. |

Alert Severity	<p>The current alert severity. Values range from 0 - 2, as indicated in the color gradient  bar, where 2 is the highest Alert Severity:</p> <ul style="list-style-type: none"> 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	<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 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.</p>
Received	<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 0 to the alarm threshold specified for the MySQLBytesReceivedHigh alert. The middle value in the gradient bar indicates the average count.</p>
Sent	<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 0 to the alarm threshold specified for the MySQLBytesSentHigh alert. The middle value in the gradient bar indicates the average count.</p>

All MySQL Instances

Investigate detailed utilization metrics for all MySQL instances. This display provides a tabular view of the performance metrics shown in the “[Instances Heatmap](#)” (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](#)” 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.

The screenshot shows the 'MySQL All Instances Table' interface. At the top, there is a title bar with a dropdown menu and a status bar indicating '11-Apr-2019 15:33' and 'No Alerts' with a green checkmark. Below the title bar, there is a search bar labeled 'Servers: 1'. The main table has the following columns: Server, Expired, Alert Level, Alert Count, Bytes Received, Bytes Sent, Connections, Delayed Writes, Queries, and Slow Query. The data row for 'MYSQLDEV' shows an Alert Level of '✓', Bytes Received of 332,013, Bytes Sent of 342,275, Connections of 9,897, Delayed Writes of 0, and Queries of 1,420.

Server	Expired	Alert Level	Alert Count	Bytes Received	Bytes Sent	Connections	Delayed Writes	Queries	Slow Query
MYSQLDEV		✓		332,013	342,275	9,897	0	1,420	

Single MySQL Instance View

Displays in this View are:

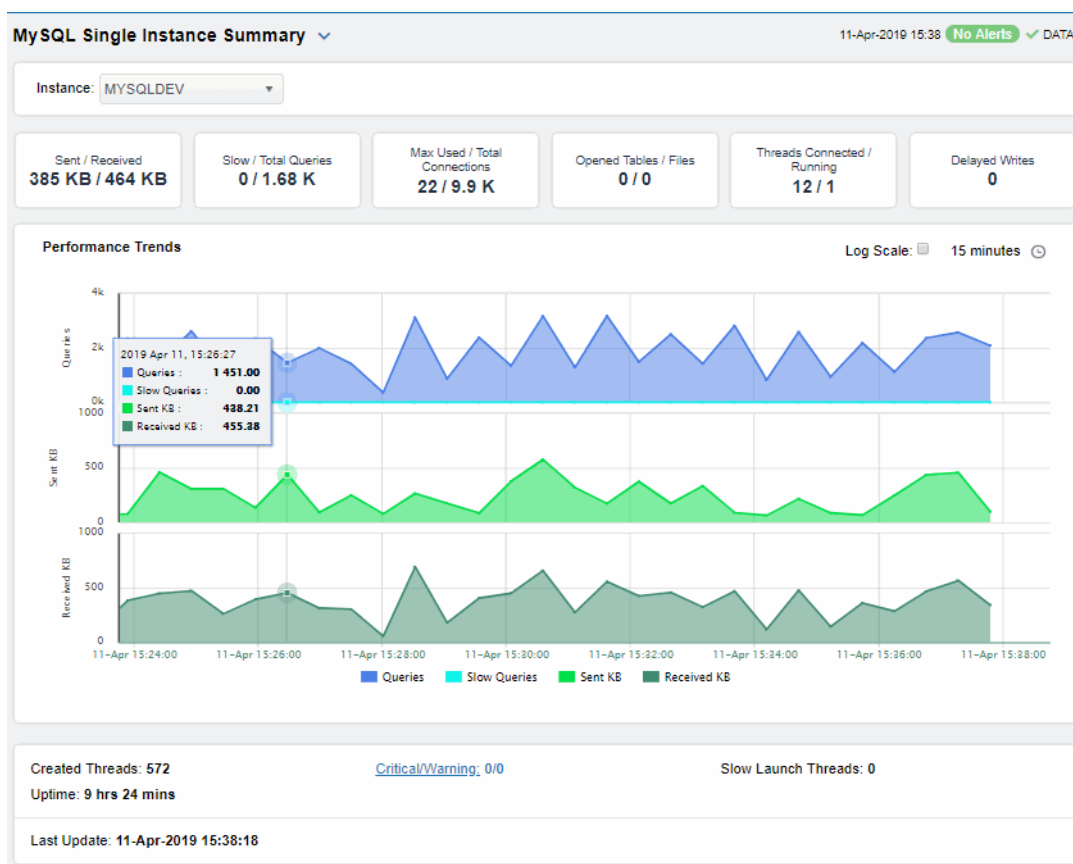
- “[Single MySQL Instance Summary](#)”: Displays performance, processing, alerts, memory, and trend data for a particular database server.
- “[Instance Properties](#)”: Displays the values of properties on servers.
- “[Instance Operations](#)”: Trend graph that traces server queries, slow queries, KB sent and KB received.
- “[Instance User Tables](#)”: A tabular view of cache tables performance and utilization metrics.

Single MySQL Instance Summary

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"](#) 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

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: MYSQLDEV

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

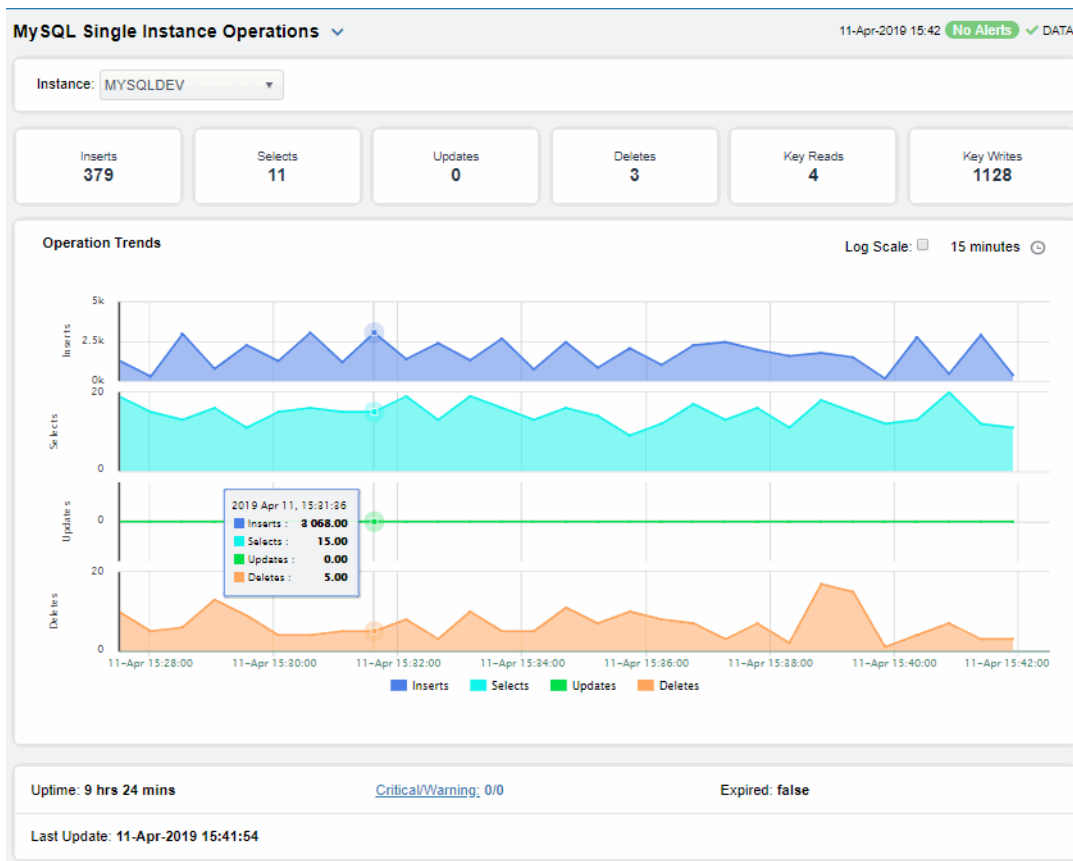
Page 1 of 7 1 - 40 of 259 items

Instance Operations

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”](#) 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

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	888,664	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
bw6_activity_totals	20,718,016	4,355,072	226,128	0	25,073,088	MyISAM	rtvhistory
bw6_appnodes	2,597,056	764,928	39,409	0	3,361,984	MyISAM	rtvhistory
bw6_process_totals	7,650,588	1,859,584	94,395	0	9,510,172	MyISAM	rtvhistory
bw6_process_totals_app	777,800	216,064	10,979	0	993,864	MyISAM	rtvhistory
bw6_process_totals_appnode	4,415,924	1,270,784	65,919	0	5,686,708	MyISAM	rtvhistory
bw6_process_totals_appslice	5,211,584	1,274,880	65,961	0	6,486,464	MyISAM	rtvhistory
bw6_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	202,651,476	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

The following can be found under **Components** tab > **Databases** > **MS SQL Server**:

- ["MS SQL Overview"](#):
- ["MS 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.
- ["MS SQL Server View"](#): The displays in this View allow you to view the metrics for a particular SQL database server.

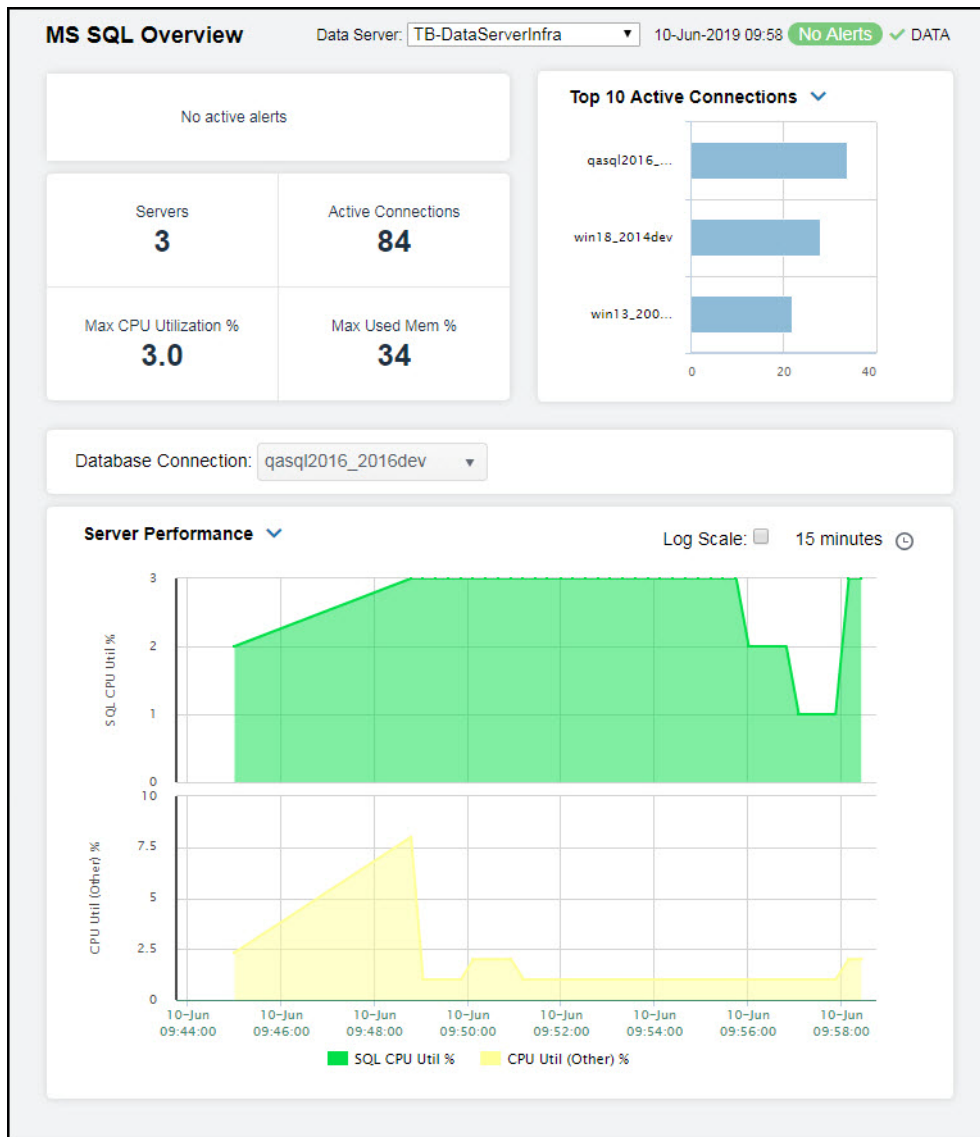
MS SQL Overview

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

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](#)” 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](#)” display, which provides a heatmap view of all MS SQL servers and their associated metrics.

MS SQL Servers Table

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](#)” 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-SQL2016MSSQLSD	✔			3.0	1.0	3.0
WIN-9J6VR0M8JK3MSE	✔			1.0	1.0	3.0
WIN13\SQLEXPRESS	✔			1.0	2.0	1.0

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.

Fields and Data

Servers The number of servers found and listed in the table.

Table

Server Name The name of the server.

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.

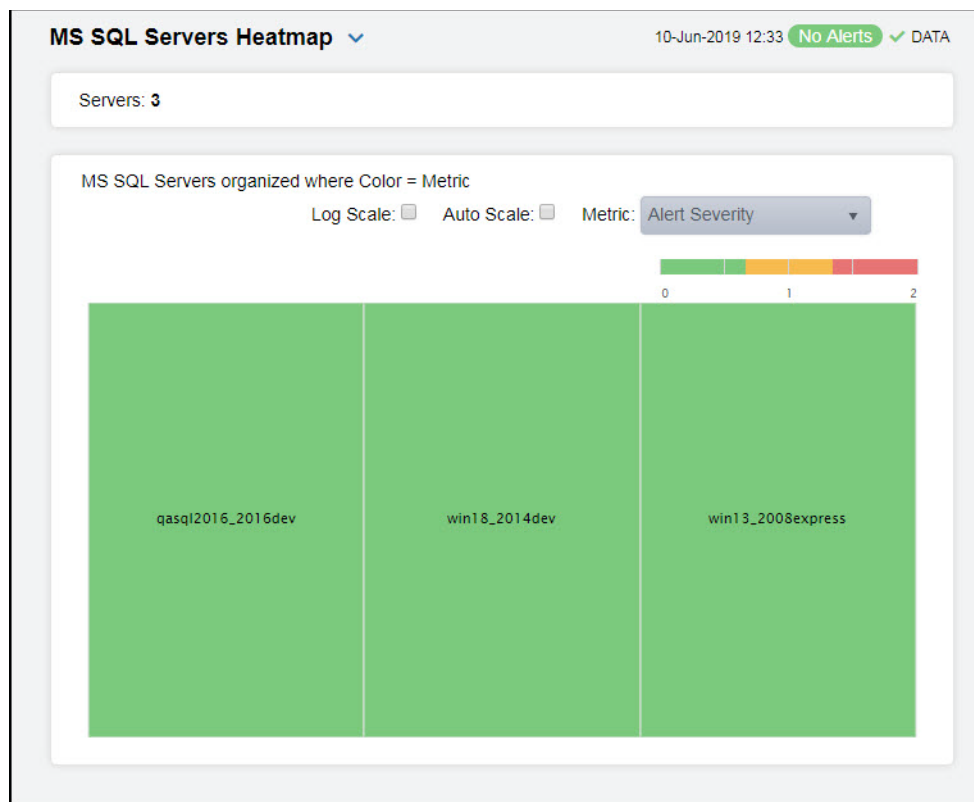
Alert Count	The total number of alerts for the host.
SQL CPU Util %	The CPU utilization percentage related to SQL.*
CPU Util (Other) %	The CPU utilization percentage for processes other than SQL.*
Used Mem %	The percentage of memory used on the server.*
Remaining Mem %	The percentage of memory remaining on the server.*
Current IO Reads	The increase in the amount of input/output reads operations (from the previous polling period to the current polling period).
Current IO Writes	The increase in the amount of input/output write operations (from the previous polling period to the current polling period).
Current IO Errors	The increase in the amount of input/output errors (from the previous polling period to the current polling period).
Active Connections	The number of currently active connections.
IO Busy ms	The time, in milliseconds, that the system has been busy due to Input/Output operations.*
Total DB Size MB	The size of the database, in megabytes.*
Current Data Sent KB	The increase in the amount of data being sent (from the previous polling period to the current polling period), in kilobytes.
Current Data Rcvd KB	The increase in the amount of data being received (from the previous polling period to the current polling period), in kilobytes.
Current Packet Errors	The increase in the amount of packet errors (from the previous polling period to the current polling period).
Product Level	The product level of the server.*
Instance	The name of the instance.
Server Edition	The version of the server.*
Product Version	The product's version number.*
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.

MS SQL Servers Heatmap

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”](#) display.



Fields and Data:

- Servers** The number of servers found and listed 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.

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 Scale is not selected.
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 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 " 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: <ul style="list-style-type: none">  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 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.

MS SQL Server View

These displays provide detailed data for a particular server. Clicking **MS SQL Server** from the left/navigation menu opens the "[MS SQL Server Summary](#)" display, which displays performance, processing, alerts, memory, and trend data for a particular database server. The options available under **MS SQL Server** are:

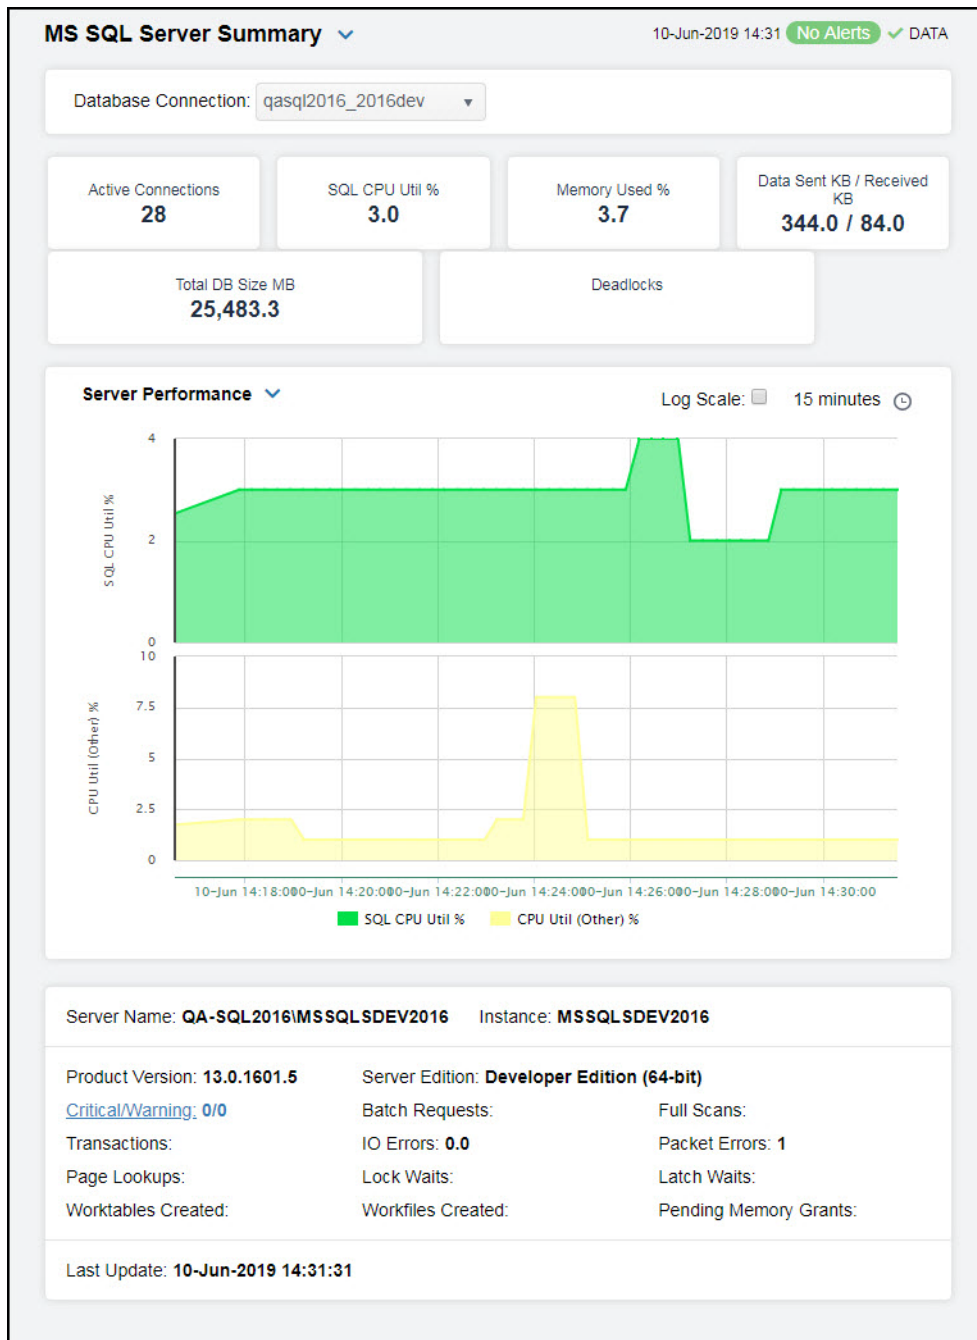
- **Instance Details:** Opens the "[MS SQL Server Performance](#)" display, which displays various database details as well as trending data for the page life expectancy.
- **Wait Stats:** Opens the "[MS SQL Server Waits Table](#)" display, which displays server wait time details in a table format for a particular database server.
- **Table Sizes:** Opens the "[MS SQL Server DB Table Sizes](#)" display, which displays database and table sizes for a particular database server.

MS SQL Server Summary

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](#)" 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.



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.

Filter By:

Database Connection Select the database for which you want to show data in the display.

Fields and Data:

Active Connections The number of active connections on the server.*

SQL CPU Util % The percentage of CPU used on SQL processing of the instance.*

Memory Used % The percentage of memory utilization of the database instance.

Data Sent KB/ Received KB The number of kilobytes sent and received.*

Total DB Size MB Total size of the database instance, in megabytes

Deadlocks The number of deadlocks on the database instance.

Trend Graphs**Server Performance**

SQL CPU Util % -- Traces the percentage of CPU used on SQL processing of the instance.

CPU Util (Other) % -- Traces the percentage of CPU used by other processes on the instance.

Server Throughput

Current Data Sent KB -- Traces the amount of kilobytes being sent.

Current Data Received KB -- Traces the amount of kilobytes being received.

Current IO Reads -- Traces the number of input/output reads.

Current IO Writes -- Traces the number of input/output writes.

Server Operations

Active Connections -- Traces the number of active connections.

Current Batch Requests -- Traces the number of batch requests since the last data update.


Current Latch Waits -- Traces the number of latch waits since the last data update.

Current Lock Waits -- Traces the number of lock waits since the last data update.

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 Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



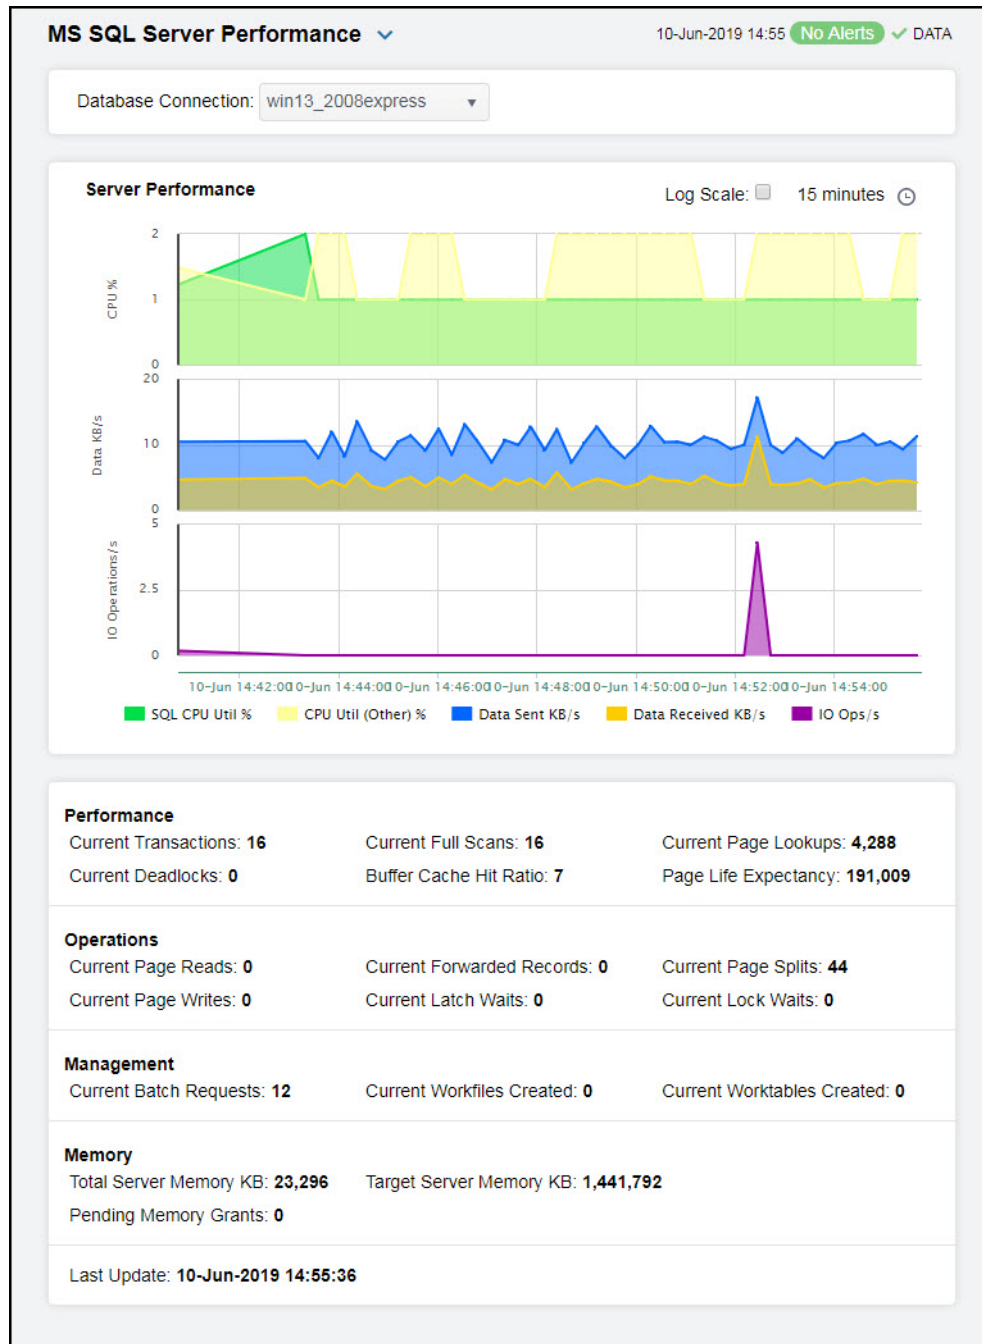
To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

Server Name	The name of the server.
Instance	The name of the instance.
Product Version	The MS SQL Server's version number.
Critical / Warning	The number of critical and warning alerts.
Transactions	The number of transactions.*
Page Lookups	The number of page lookups.
Worktables Created	The number of worktables created.
Server Edition	The SQL Server's edition.*
Batch Requests	The number of batch requests on the server.
IO Errors	The number of input/output operation errors on the server.
Lock Waits	The number of lock waits.
Workfiles Created	The number of workfiles created.
Full Scans	The number of full scans executed.
Packet Errors	The number of packet errors.
Latch Waits	The number of latch waits.

Pending Memory Grants	The number of pending memory grants.
Last Update.	The date and time of the last data update.

MS SQL Server Performance

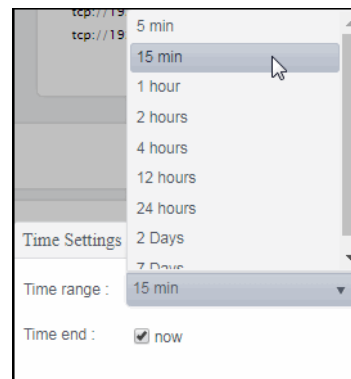
Clicking **Instance Details** in the left/navigation menu opens the **MS SQL Server Performance** display, which allows you to view various database details as well as trending data for the page life expectancy.



Filter By:

The display might include these filtering options:

- | | |
|---------------------------------|---|
| Database Connection | Select the database for which you want to show data in the display. |
| Performance Trends Graph | Traces the following: <ul style="list-style-type: none"> SQL CPU Util % -- Traces percentage of CPU used by the MS SQL Server. CPU Util (Other) % -- Traces the percentage of CPU used in other operations by the MS SQL Server. Data Sent KB/s -- Traces the rate of data sent, in kilobytes per second. Data Received KB/s -- Traces the rate of data received, in kilobytes per second. IO Operations/s -- Traces the rate of input/output 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. |
| Time Settings | Select a time range from the drop down menu varying from 5 Minutes to Last 7 Days . By default, the time range end point is the current time. |



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

Performance

Current Transactions	The number of transactions since the last data update.*
Current Deadlocks	The number of deadlocks since the last data update.*
Current Full Scans	The number of full scans since the last data update.*
Buffer Cache Hit Ratio	The current buffer cache hit ratio, which is the total number of cache hits divided by the total number of cache lookups.*
Current Page Lookups	The number of page lookups since the last data update.*
Page Life Expectancy	The average number of seconds a page stays in the cache.*

Operations

Current Page Reads	The number of page reads since the last data update.*
Current Page Writes	The number of page writes since the last data update.*
Current Forwarded Records	The number of forwarded records since the last data update.*
Current Latch Waits	The number of latch waits since the last data update.*
Current Page Splits	The number of page splits since the last data update.*
Current Lock Waits	The number of lock waits since the last data update.*

Management

Current Batch Requests	The number of batch requests since the last data update.*
Current Workfiles Created	The number of work files created since the last data update.*
Current Worktables Created	The number of worktables created since the last data update.*

Memory

Total Server Memory KB	The total amount of memory the server has committed using the memory manager, in kilobytes.*
Target Server Memory KB	The defined target server memory, which is the ideal amount of memory the server can consume, in kilobytes.*

Pending Memory Grants

The current number of processes waiting for a workspace memory grant.*

Last Update

The date and time of the last data update.

MS SQL Server Waits Table

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.

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

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.

Filter By:

The display includes these filtering options:

Database Connection

Select the database for which you want to show data in the display.

Wait Types Sorted By Non-Zero Percentage Per Wait Category Table

Wait Category	The name of the wait category.*
Wait Type	The name of the wait type.*
Wait s	The average length of the wait time, in seconds.*
Signal s	When the thread is marked as runnable, this field displays the wait time, in seconds, that it takes to get into the running state.*
Resource s	The length of time the thread spent in a suspended state waiting to acquire a resource, in seconds.*
Percentage	The percentage of time the thread spent in a wait state for this wait type.*
Wait Count	The number of lock requests that required the caller to wait.*
Avg Wait s	The average wait time, in seconds.*
Avg Signal s	The average wait signal time, in seconds.*
Avg Resource Time s	The average length of time taken to acquire a resource, in seconds.*
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.
Time Stamp	The date and time the row data was last updated.

MS SQL Server DB Table Sizes

Clicking **Table Sizes** in the left/navigation menu opens the **MS SQL Server DB 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	

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.

Filter By:

Database Connection Select the database for which you want to show data in the display.

Database Sizes Table

Database The name of the database.*

Log Size MB The size of the log, in megabytes.*

Row Size MB The row size, in megabytes.*

State The current state of the database.*

Total Size MB The total size of the database, in megabytes.*

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.
----------------	--

Table Sizes

Database	The name of the database.*
Schema	The name of the schema.*
Table	The name of the table.*
Row Count	The number of rows in the table.*
Total Space (MB)	The total space, in megabytes, available in the table.*
Used Space (MB)	The total number of used megabytes in the table.*
Unused Space (MB)	The total number of unused megabytes in the table.*
Creation Date	The date and time the table was created.*

Node.js

This section describes 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](#)”: The displays in this View present detailed data for all node instances or for a particular node instance.
- “[Node Requests 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.
- “[Node Processes 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.

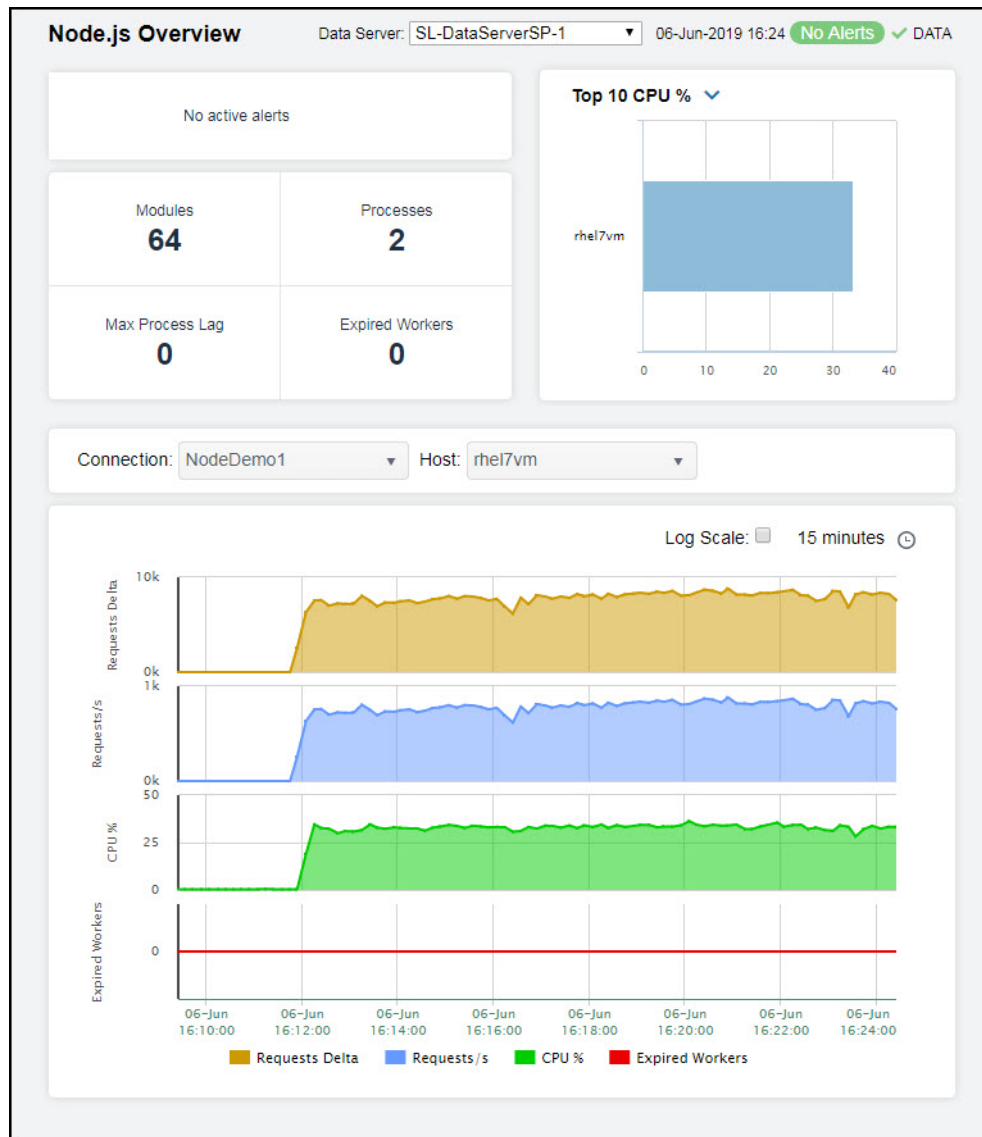
Node.js Overview

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

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](#)” 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](#)” display, which provides a way to view trending data for individual node processes.

Node/Master Info 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 double-clicking a row to view details for the selected node in the “[Node Master Summary](#)” 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

Fields and Data:

- Connection** Select the name of the connection containing the node instances for which you want to view data.
- Masters** The total number of node instances being monitored based on your search criteria.

Table:

Each row in the table is a different message router.

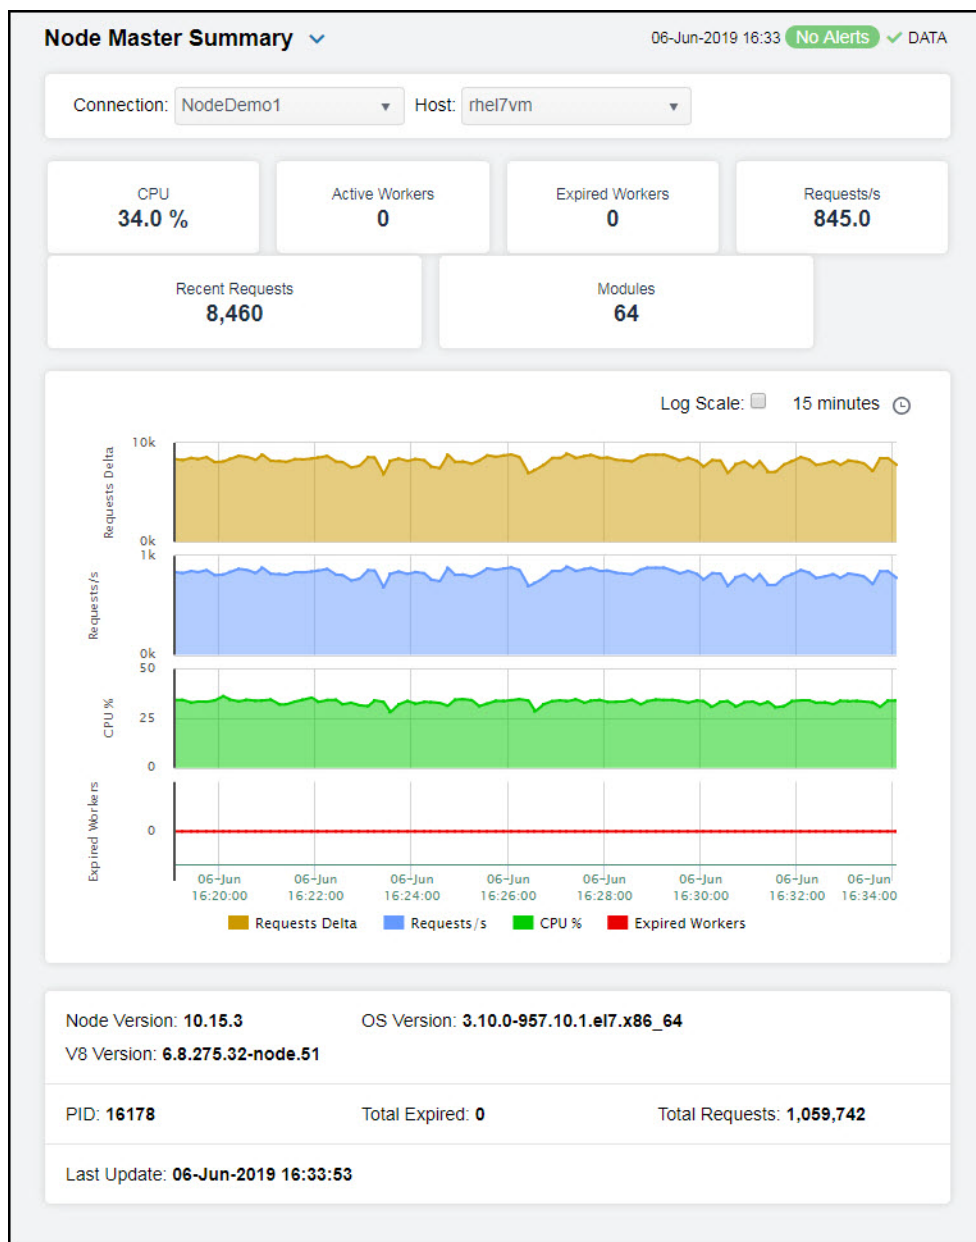
- Connection** The name of the connection.
- Host** The name of the host.
- Alert Level** The current alert severity.
- Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
 - Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
 - Green indicates that no metrics have exceeded their alert thresholds.
- Alert Count** The total number of current alerts.

Process ID	The process id for the node instance.
Up Time	The amount of time the process has been running.
CPU %	The percentage of CPU used for the process.
Request Count	The total number of requests on the host.
Requests/s	The average number of requests per second on the host.
Requests Delta	The total number of requests since the last data update.
Request Mean Rate	The average number of requests for the server since monitoring was started.
Requests 1 Min Rate	The average number of requests for the last minute.
Request 5 Min Rate	The average number of requests for the last 5 minutes.
Request 15 min Rate	The average number of requests for the last 15 minutes.
Expired Workers	The number of expired workers on the host since the last data update.
Arch	The CPU architecture of the operating system on the server. Possible values are x64 , arm , and ia32 .
ares	The current version of C-ares running on the host.
HTTP Parser	The current version of the http parser running on the host.
ICU	The current version of ICU running on the host.
Modules	This number of modules found on the host.
Node Version	The version of node.js running on the host.
OpenSSL	The current version of OpenSSL running on the host.
Platform	The operating system's platform. Possible values, among others, are: darwin , linux , sunos , or win32 .
OS Version	The operating system's release number.
Type	The name of the operating system. Possible values, among others, are Linux on Linux, Darwin on OS X, and Windows_NT on Windows.
UV	The current version of uv running on the host.
V8 Version	The current version of v8 running on the host.
ZLib	The current version of ZLib running on 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 > 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.
Time Stamp	The date and time the row data was last updated.

Node Master Summary

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”](#) 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.



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:

CPU The percentage of memory used on the CPU.

Active Workers The current number of active workers on the host.

Expired Workers The total number of expired workers on the host.

Requests/s The average number of requests per second on the host.

Recent Requests The total number of requests since the last data update.

Modules This number of modules found on the host.

Trend Graph

Traces the following:

Request s Delta -- traces the number of requests on the host.

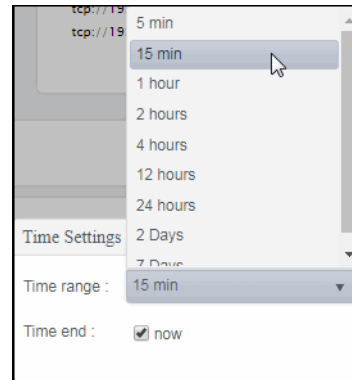
Requests/s -- traces the number of requests per second on the host.

CPU % -- traces the percentage of CPU being used on the host.

Expired Workers -- traces the number of expired workers on the host.

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 Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

Node Version	The version of node.js running on the host.
V8 Version	The current version of v8 running on the host.
OS Version	The operating system's version number.
PID	The process id for the node instance.
Total Expired	The total number of expired workers on the host.
Total Requests	The total number of requests on the host.
Last Update	The date and time of the last data update.

Node Requests 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. Clicking **Node Requests** from the left/navigation menu opens the “[Node Requests Table](#)” 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](#)” 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

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. 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	/

Filter By:

- 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:

Request URLs: The total number of request URLs (rows) in the table.

Table:

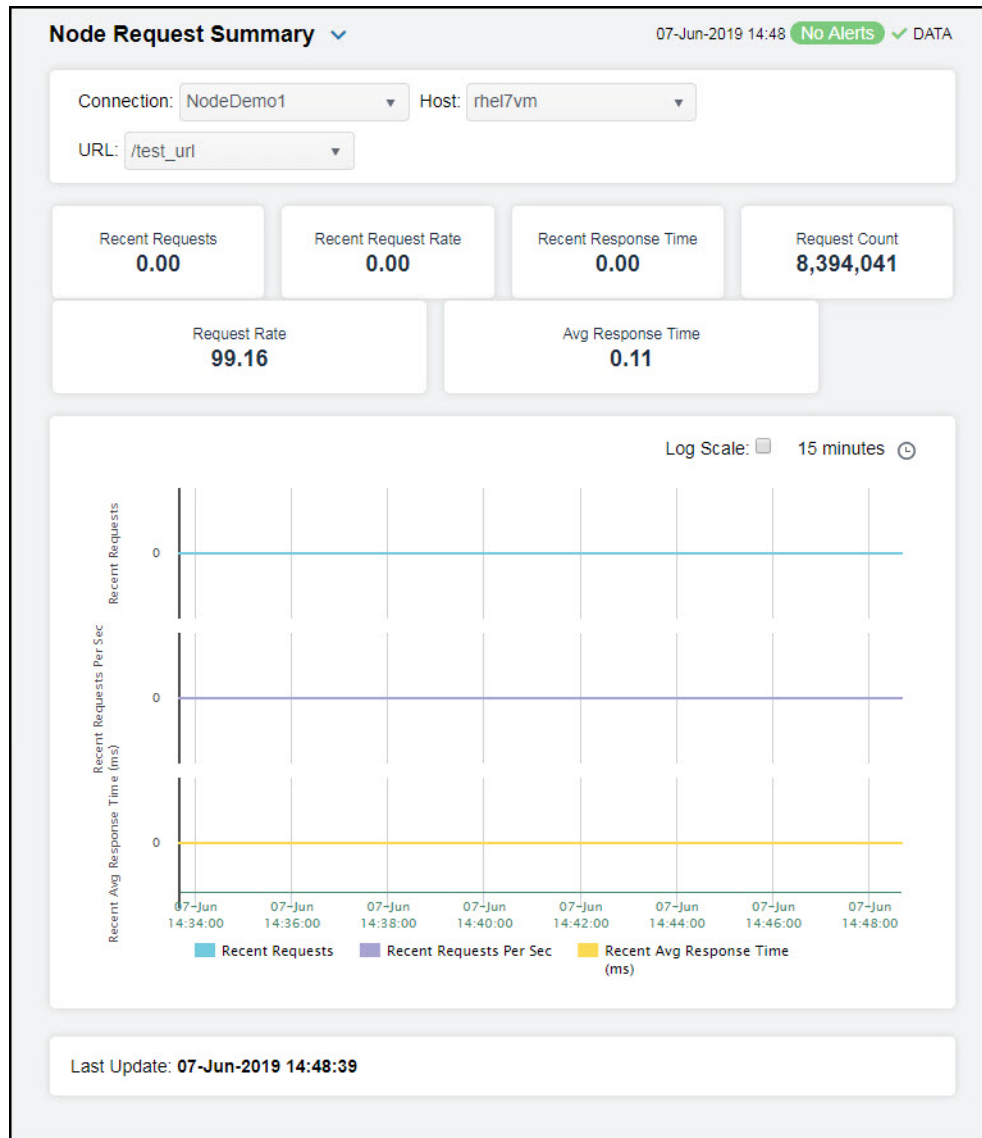
Column values describe the node and its associated requests.

Connection	The name of the connection
Host	The name of the host.
Request URL	The URL from which the requests originated.
Request Count	The total number of requests.
Requests / Sec	The rate of requests since the server was started.
Avg Resp Time (ms)	The average response time (in milliseconds) since the server was started.
Recent Requests	The total number of requests based on the last query interval.
Recent Requests Per Sec	The rate of recent requests based on the last query interval.
Recent Avg Response Time (ms)	The average response time (in milliseconds) based on the last query interval.
Time Stamp	The date and time the row data was last updated.

Node Request Summary

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.



Filter By:

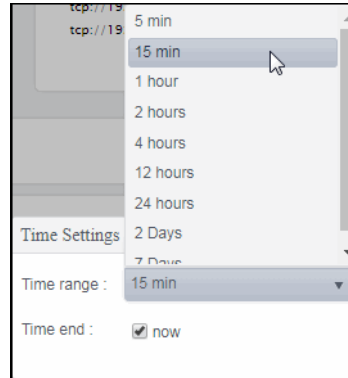
- Connection** Select the connection for which you want to show data in the display.
- Host** Select the host for which you want to show data in the display.
- URL** Select the URL for which you want to view data.

Fields and Data:

Recent Requests	The total number of requests since the last data update.
Recent Request Rate	The rate of requests since the last data update.
Recent Response Time	The average response time since the last data update.
Request Count	The total number of requests.
Request Rate	The rate of requests.
Avg Response Time	The average response time (in milliseconds).
Performance Trends Graph	<p>Traces the following:</p> <ul style="list-style-type: none"> Recent Requests -- traces the total number of requests since the last data update. Recent Requests Per Sec -- traces the rate of requests since the last data update. Recent Avg Resp Time (ms) -- traces the average response time, in milliseconds, since the last data update. <p>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.</p>

Time Settings

Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

Last Update The date and time of the last data update.

Node Processes 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. Clicking **Node Processes** from the left/navigation menu opens the ["Node Processes Table"](#) 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"](#) 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"](#) display, which allows you to view current and trending data for a single node process for a particular connection/host combination.

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 double-clicking on a row in the resulting 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.

Show Masters Only Select this check box to view only masters in the table.

Processes Lists the total number of processes (rows) found using the search parameters.

Table

Connection The name of the connection.

Host The name of the host.

Master / Worker Displays whether the process is the Master process or, if the application is clustered, the worker ID.

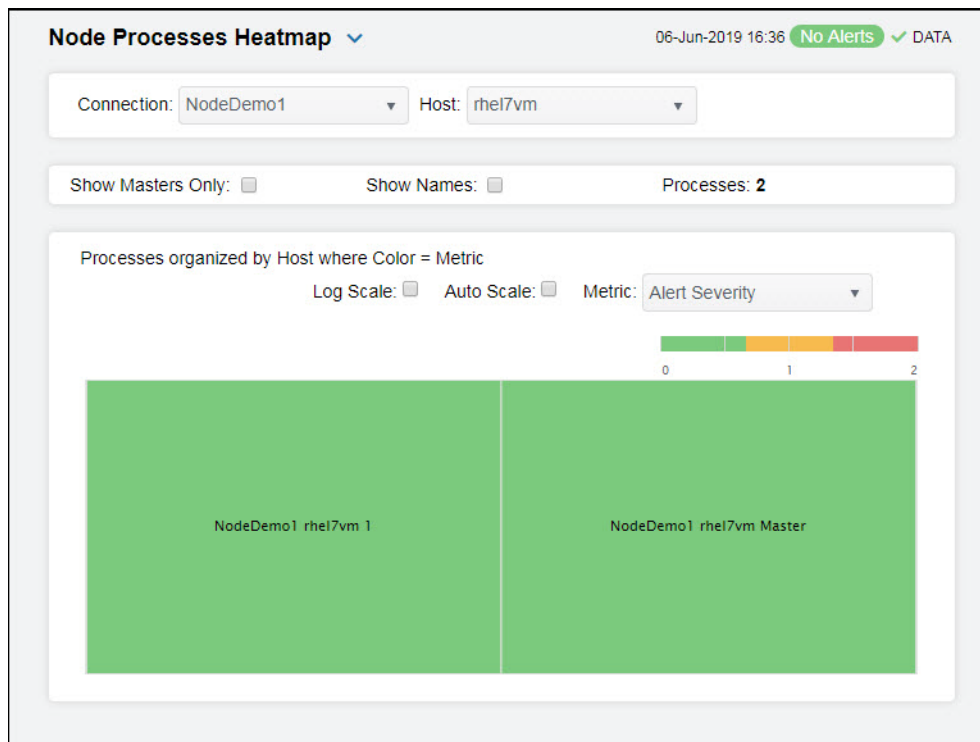
Alert Level	<p>The current alert status.</p> <ul style="list-style-type: none"> ● Red indicates that one or more metrics exceeded their ALARM LEVEL threshold. ● Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold. ● Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	Total number of alerts for the process.
Up Time	Lists the amount of time the process has been up and running.
CPU %	A decimal percentage describing how much the process utilizes the CPU.
Process ID	The process ID.
Memory Used	The used memory as a fraction of total system memory, in kilobytes.
Memory Used %	The percentage of total available memory used.
Mem.rss	The Resident Set Size, which is the portion of memory held in RAM (as opposed to swap or disk), in kilobytes.
Heap Total	The total amount of heap memory from which newly created objects will originate, in kilobytes.
Heap Free	The amount of memory remaining from which newly created objects will originate, in kilobytes.
Heap Used	The heap memory currently in use, in kilobytes.
Heap Used %	The percentage of heap memory currently being used.
Heap Available	The v8 engine's total_available_size value, in kilobytes.
Heap Limit	The v8 engine's heap_size_limit value, in kilobytes.
Heap Total Executable	The v8 engine's total_heap_size_executable value, in kilobytes.
Latency p100	The number of microseconds that 100 percent of events were late in the previous 4 seconds.
Latency p99	The number of microseconds that 99 percent of events were late in the previous 4 seconds.
Latency p95	The number of microseconds that 95 percent of events were late in the previous 4 seconds.
Latency p90	The number of microseconds that 90 percent of events were late in the previous 4 seconds.
Latency p50	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.

- 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.
- Time Stamp** The date and time the row data was last updated.

Node Processes Heatmap

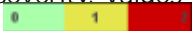

Clicking **Node Processes Heatmap** in the left/navigation menu opens the **Node Processes Heatmap** display, which provides an easy-to-view interface that allows you to quickly identify the current status of each of your node processes for each available metric. You can view the node processes in the 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”](#) display, which allows you to see additional details for the selected node process.



Fields and Data:

- Connection** Select the connection for which you want to show data in the display.

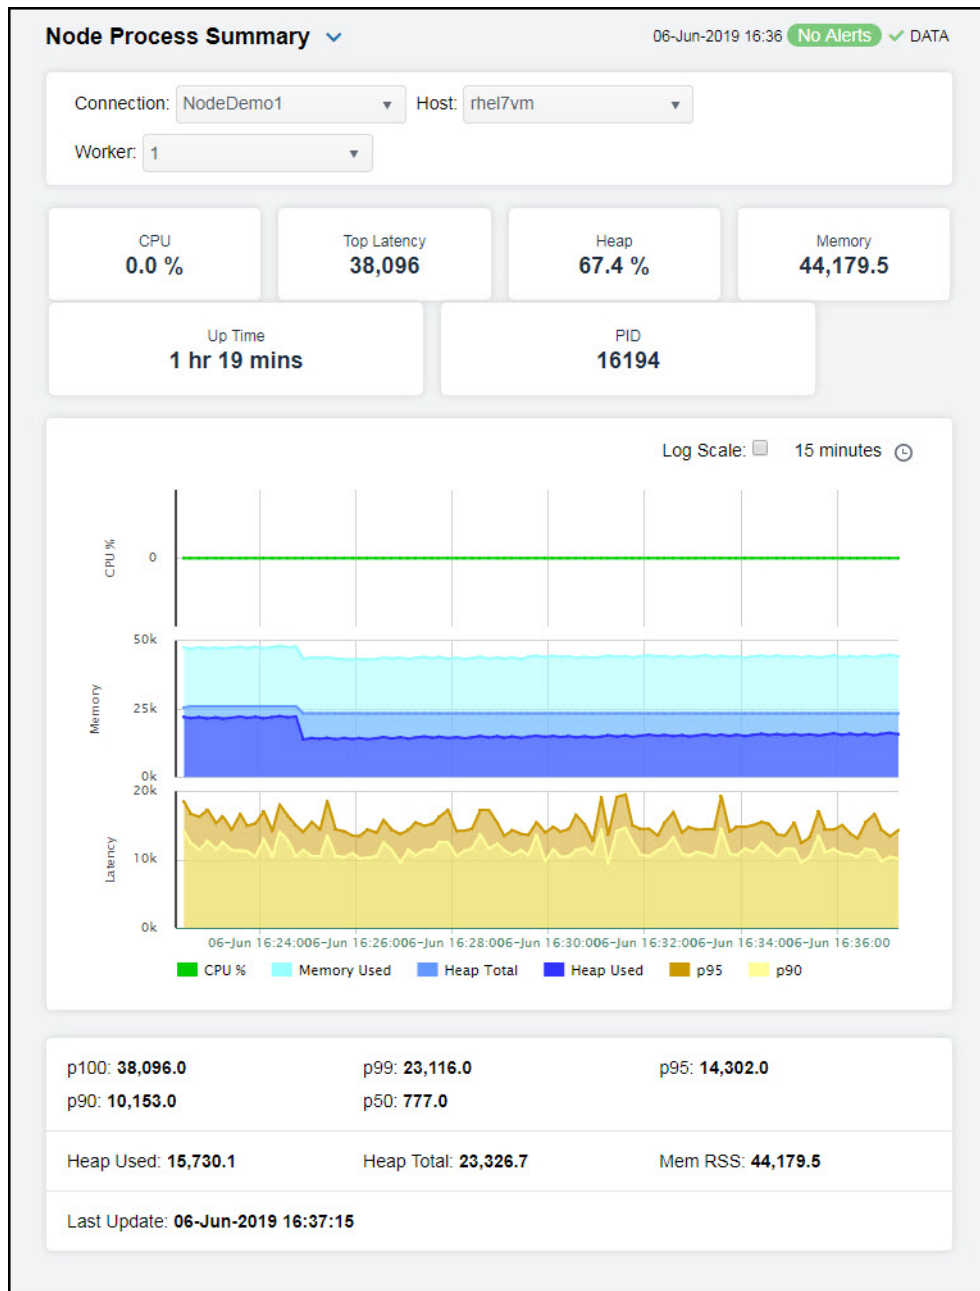
Host	Select the host for which you want to show data in the display.
Show Masters Only	Select this check box to view only masters in the heatmap.
Show Names	Select this check box to display the names of the processes at the top of each rectangle in the heatmap.
Processes	Lists the total number of processes (rows) found using the search parameters.
Heatmap	
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.
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 Scale is not selected.
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 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 " display for a detailed view of metrics for that particular process.
Alert Severity	<p>The current alert severity. Values range from 0 - 2, as indicated in the color gradient  bar, where 2 is the highest Alert Severity:</p> <ul style="list-style-type: none"> ● 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 %	<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 NodeProcessCpuUsageHigh. The middle value in the gradient bar indicates the middle value of the range.</p> <p>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.</p>
Memory Used %	<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 NodeProcessMemUsageHigh. The middle value in the gradient bar indicates the middle value of the range.</p> <p>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.</p>

Node Process Summary

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.


**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.
- Worker** Select the name of the worker to view. You can select from **Master** or any of the worker processes created by the Master. Worker processes are defined by numbers: **1** for the first worker process created by the **Master**, **2** for the second worker process created by the **Master**, and so on.

Fields and Data

CPU	A decimal percentage describing how much the process utilizes the CPU.
Top Latency	The number of microseconds that 100 percent of events were late in the previous 4 seconds.
Heap	The decimal percentage of utilized heap space.
Memory	The total amount of memory from which newly created objects can originate, in kilobytes.
Up Time	Lists the amount of time the process has been up and running.
PID	The process ID.
Performance Trends Graph	Traces the following: CPU % -- traces the CPU utilization percentage. Memory 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.

Time Settings Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

p100	The number of microseconds that 100 percent of events were late in the previous 4 seconds.
p90	The number of microseconds that 90 percent of events were late in the previous 4 seconds.
p99	The number of microseconds that 99 percent of events were late in the previous 4 seconds.
p50	The number of microseconds that 50 percent of events were late in the previous 4 seconds.
p95	The number of microseconds that 95 percent of events were late in the previous 4 seconds.
Heap Used	The heap memory currently in use, in kilobytes.
Heap Total	The total amount of memory from which newly created objects can originate, in kilobytes.
Mem RSS	Resident Set Size, which is the portion of memory held in RAM (as opposed to swap or disk), in kilobytes.
Last Update	The date and time of the last data update.

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.

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 Solution Package for RTView Host Agent features an overview display, [“Hosts Overview”](#), and the following displays which can be found under **Components** tab > **Hosts**:

- [“Hosts Heatmap”](#)
- [“Hosts Table”](#)
- [“Host Summary”](#)
- [“Host Processes”](#)
- [“Host Network Interfaces”](#)
- [“Host Storages”](#)

Hosts Overview

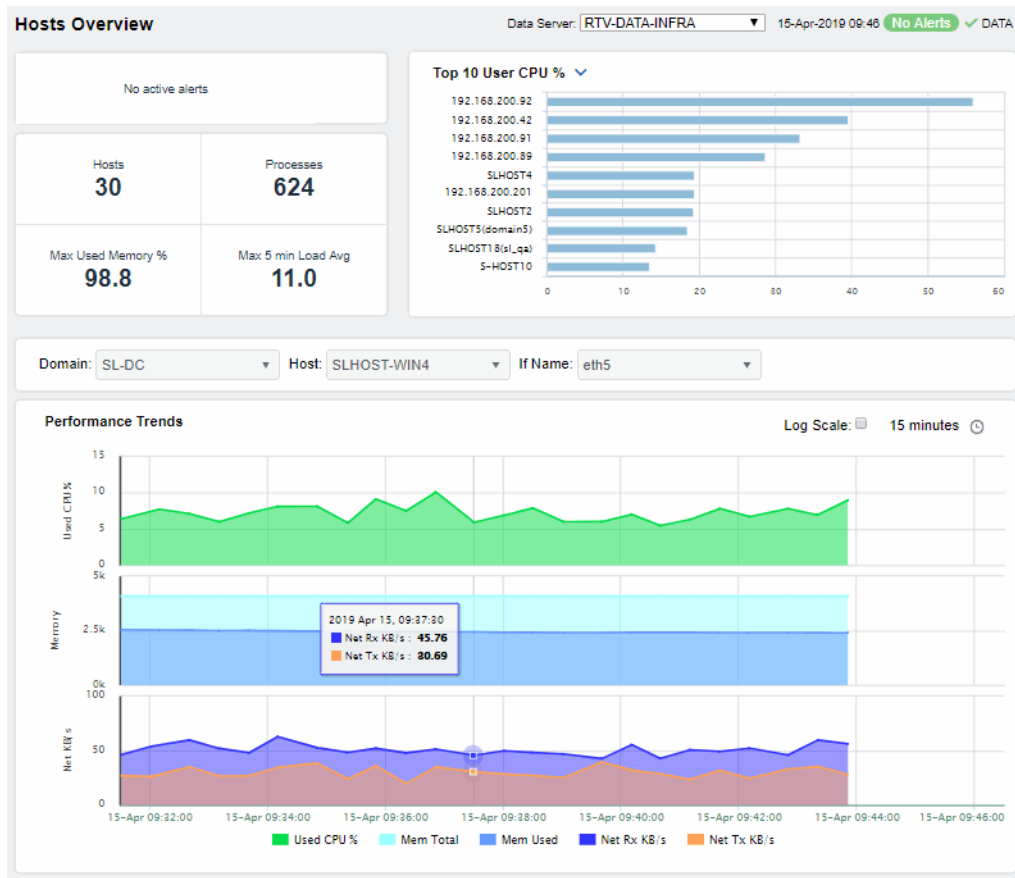
The Hosts Overview is the top-level display for the Hosts Solution Package, which provides a good starting point for immediately getting the status of all your connections on your Data Server. You can select the RTView DataServer for which you want to see data and easily view the current data for that 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”](#) 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

View the status and most critical alert states of all your hosts. Use this display to quickly identify hosts with critical alerts.

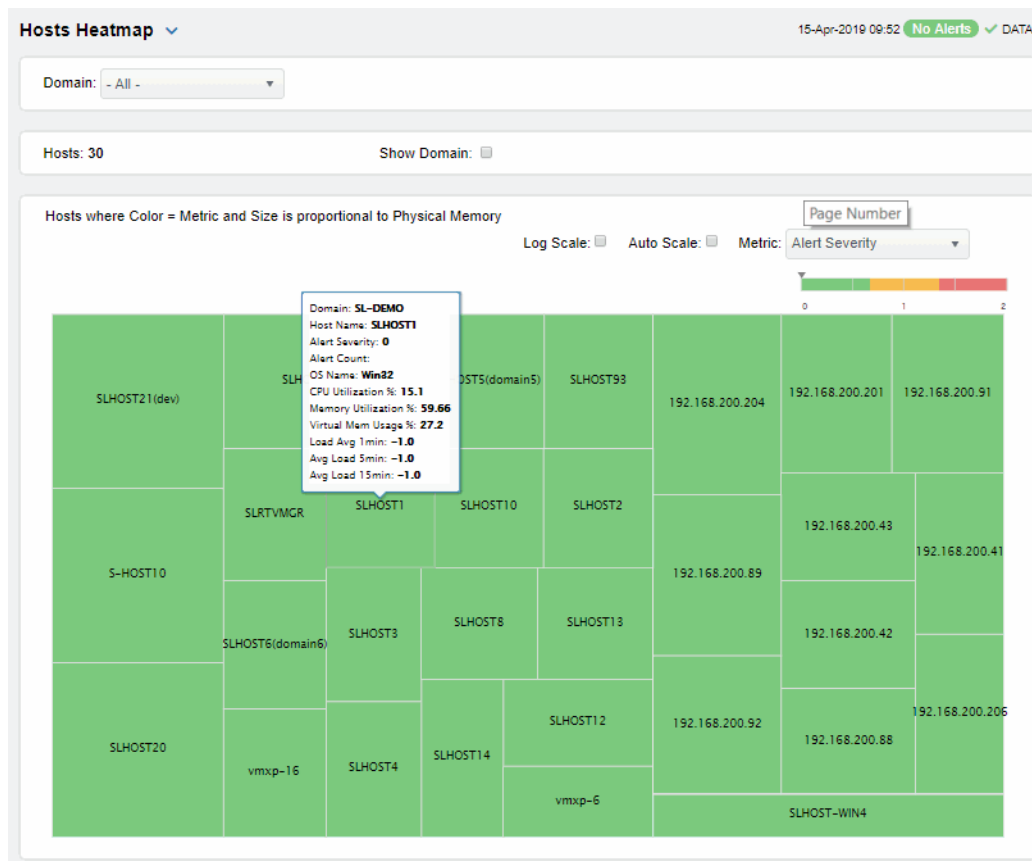
Each rectangle in the heatmap represents a host. The rectangle color indicates the most critical alert state associated with the host for the selected **Metric**. The rectangle size represents the amount of physical memory present on the host; a larger size is a larger value.

- 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” display.



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.


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 Utilization

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 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

% Memory Used 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 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

% Virtual Memory Used 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 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

Load Avg 1 Min The average number of processes running over 1 minute.

Hosts Table

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** 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

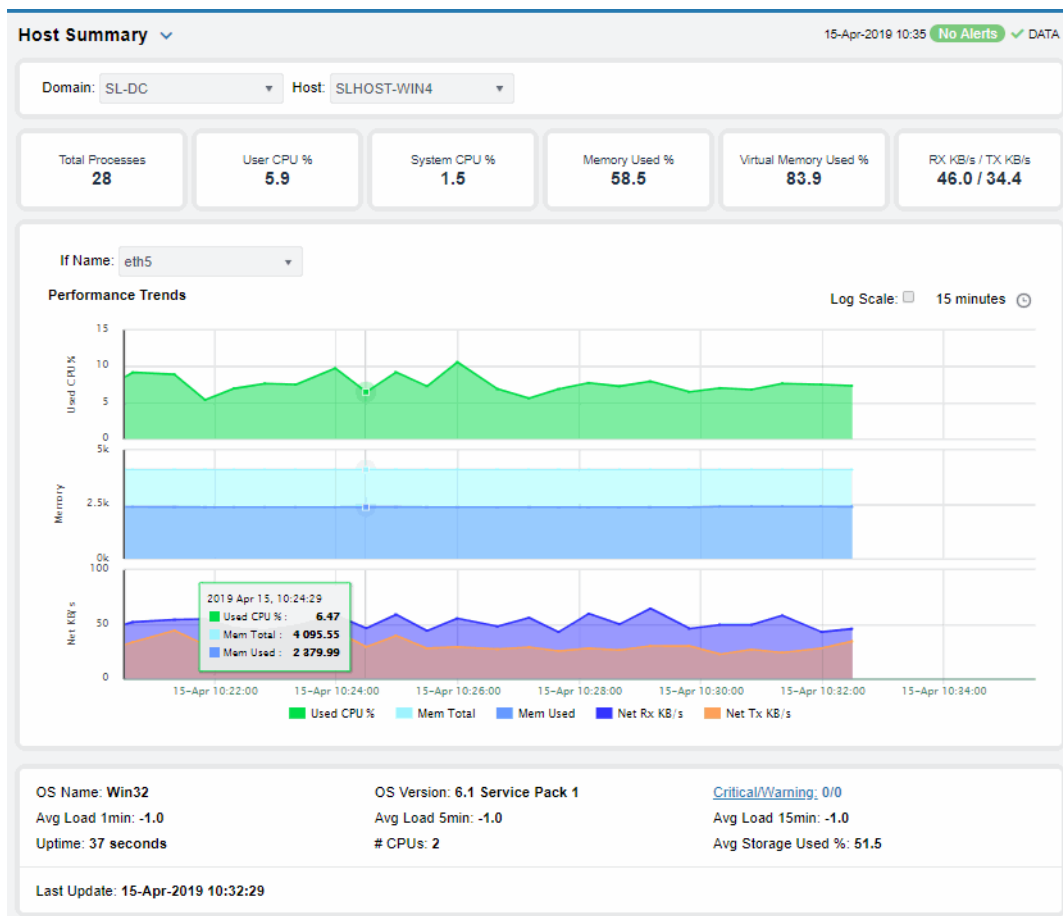
Domain	Host Name	Expired	Alert Level	Alert Count	Uptime	% CPU User	% CPU System	Idle CPU %	Used Memory
SL-DC	SLHOST-WIN4				36 seconds	7.2	1.5	91.3	2
SL-DEMO	S-HOST10				36 seconds	14.5	3.2	82.1	5
SL-DEMO	SLHOST1				15 seconds	10.9	2.9	85.6	1
SL-DEMO	SLHOST10				37 seconds	2.3	0.6	96.8	2
SL-DEMO	SLHOST12				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

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](#)” 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

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: Host:

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.gmsjiau	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.gmsjiau	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.gmsjiau	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.gmsjiau	0.0	15-Apr-2019 00:06	1,356,103,680
SL-DC	SLHOST-WIN4	🔴	3920	Administrator	java:com.sl.gmsjiau	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.gmsjiau	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.gmsjiau	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.gmsjiau	0.0	15-Apr-2019 00:06	1,356,103,680
SL-DC	SLHOST-WIN4	🔴	804	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.gmsjiau	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.gmsjiau	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

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: - All - Host: - All -

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

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

This section describes 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](#)”: 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](#)”: View current and historical data for your virtual machines.

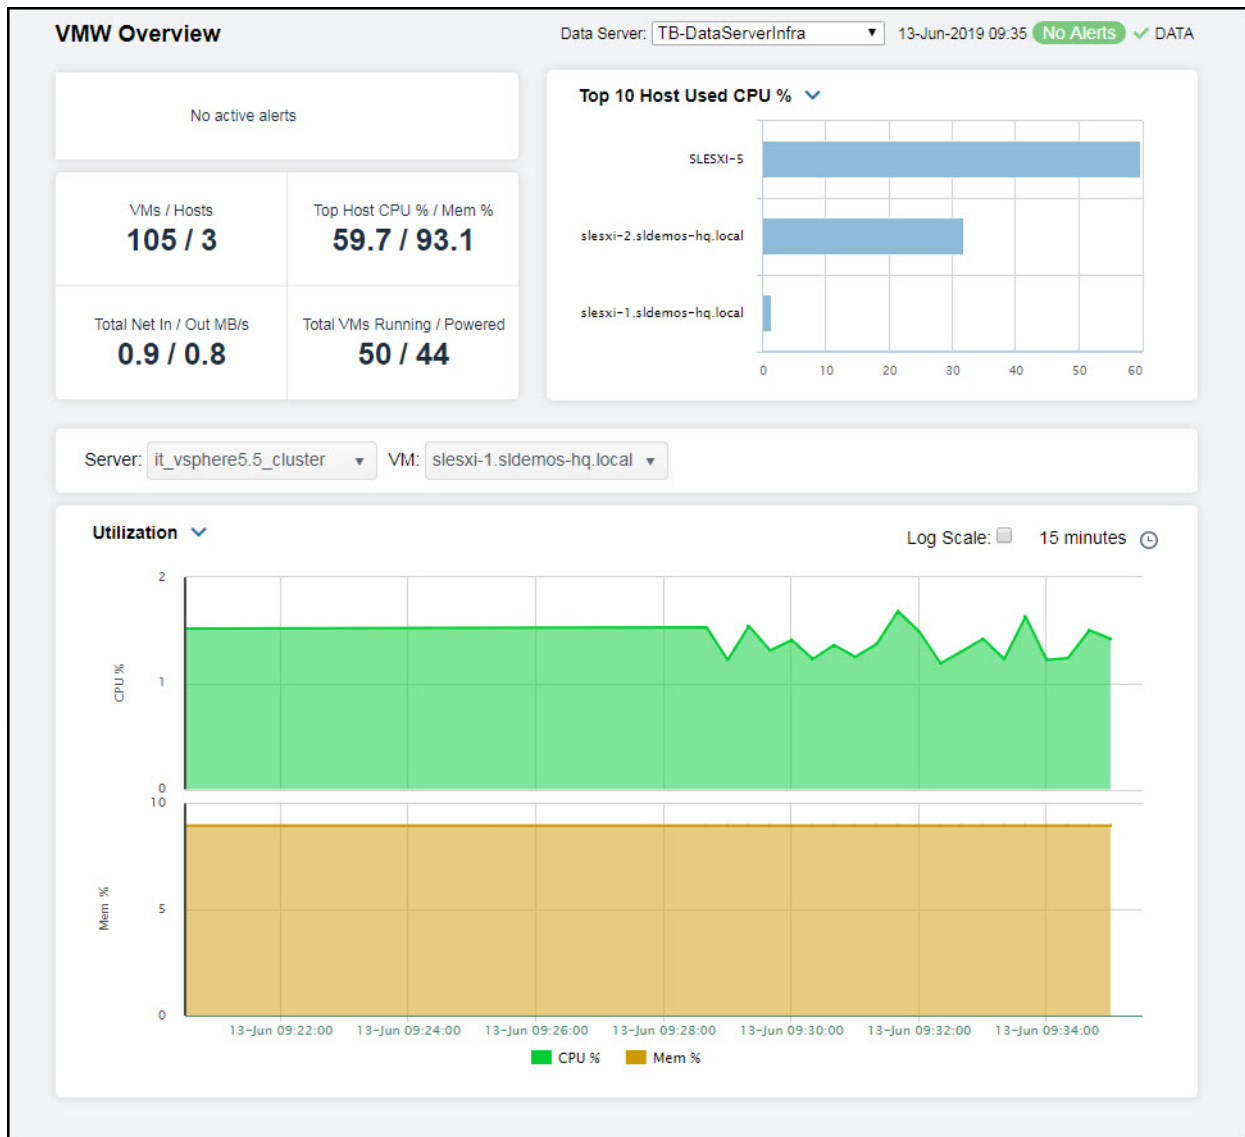
VMware Overview Display

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

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](#)” 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](#)”, which shows utilization data for a specific host running virtual machines.

VMware Hosts Table

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](#)” display and view metrics for that particular host.

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	

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:

Server Select the server for which you want to view data.

Hosts The total number of hosts of the selected cluster(s), which are listed in the table

Hosts Table

Server The name of the server.

Host Name The name of the host.

Alert Level	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.
Alert Count	The number of alerts currently on the host.
Overall 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.
Connection State	Lists the status of the connection.*
Power State	Lists whether the host is powered on or powered off.*
In Maintenance Mode	When checked, this check box signifies that the host is in maintenance mode.*
Standby Mode	The host's standby mode.*
Host Used CPU %	The percentage of CPU used by the virtual machines.*
Cores	The total number of cores on the CPU.*
CPU Threads	The total number of threads on the CPU.*
Host Memory Used %	The percentage of the host's memory currently in use.*
Memory Used	The total memory used, in megabytes, on the host.*
Memory Size	The total amount of memory, in megabytes.*
Swap Mem	The total amount of swap space used by the host, in megabytes.*
Disk Reads KB/s	The amount of data being read from the disk per second, in kilobytes.*
Disk Writes KB/s	The amount of data being written to the disk per second, in kilobytes.*
Host IN KB/s	The amount of network data being received per sec, in kilobytes.*
Host OUT KB/s	The amount of network data being transmitted per sec, in kilobytes.*
Rx Dropped Packet %	The percentage of incoming packets that were dropped.*
Tx Dropped Packets %	The percentage of outgoing packets that were dropped.*
Rx Error Packets %	The percentage of incoming packets that had errors.*

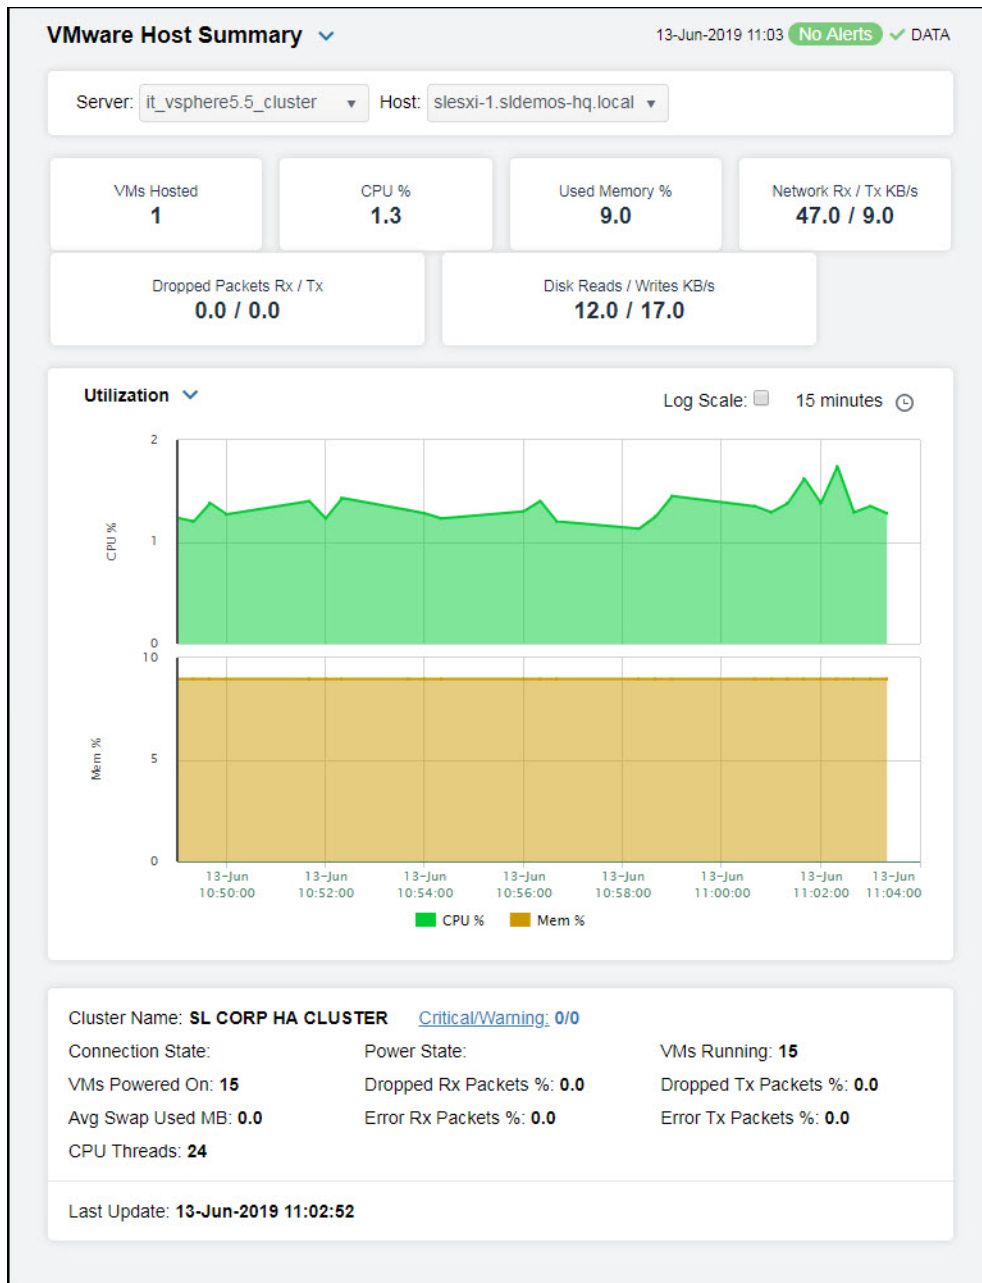
Tx Error Packets %	The percentage of outgoing packets that had errors.*
Pkts IN	The number of incoming packets.*
Pkts OUT	The number of outgoing packets.*
Pkts IN Dropped	The number of incoming packets that were dropped.*
Pkts OUT Dropped	The number of outgoing packets that were dropped.*
Pkts IN Errors	The number of incoming packets that had errors.*
Pkts OUT Errors	The number of outgoing packets that had errors.*
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.
Time Stamp	The date and time the data was last updated.

VMware Host Summary

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”](#) 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.



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:

- Server** Select the server containing the host for which you want to view data.
- Host** Select the host of the virtual machines for which you want to view data.

Fields and Data:

- VMs Hosted** The number of virtual machines on the host.*
- CPU %** The percentage of average CPU used on the host.*
- Used Memory %** The percentage of average used memory on the host.
- Network Rx/Tx KB/s** The rate of network received and transmitted data, in kilobytes per second, by this host.
- Dropped Packets Rx/Tx** The total number of dropped packets that were received and transmitted on this host.
- Disk Reads/Writes KB/s** The rate of disk reads and writes, in kilobytes per second, on this host.

Trend Graphs**Utilization**

- CPU Usage** -- Traces the percentage of CPU used.
- Memory Usage** -- Traces the amount of memory used.

Disk

- Reads KB/s** -- Traces the amount of data being read from the disk, in kilobytes per second.
- Writes KB/s** -- Traces the amount of data being written to the disk, in kilobytes per second.
- CPU %** -- Traces the percentage of CPU used.
- Mem %** -- Traces the percentage of memory used.

Network

- Net Rx KB/s:** The amount of network data received per second, in kilobytes per second.
- Net Tx KB/s:** The amount of network data transmitted per second, in kilobytes per second.
- CPU %** -- Traces the percentage of CPU used.
- Mem %** -- Traces the percentage of memory used.

Network Loss

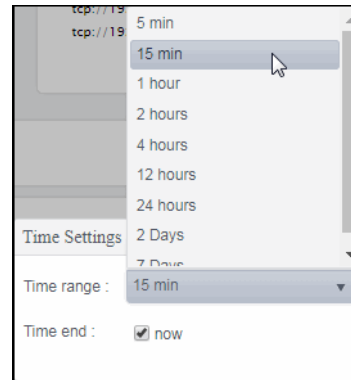
- Pkts Rx** -- Traces the number of incoming packets that were dropped.
- Pkts Tx** -- Traces the number of outgoing packets that were dropped.

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

Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

Cluster Name	The name of the cluster.
Power State	Lists whether the host is powered on or powered off.*
Dropped Rx Packets %	The percentage of incoming packets that were dropped.*
Error Rx Packets %	The percentage of incoming packets that had errors.*
Critical/Warning	The number of critical and warning alerts.
VMs Running	The number of virtual machines currently up and running on the host.*
Dropped Tx Packets %	The percentage of outgoing packets that were dropped.*
Error Tx Packets %	The percentage of outgoing packets that had errors.*
Connection State	Displays the current state of the connection for the host (connected/notConnected).*
VMs Powered On	The number of virtual machines on the host that are powered on.*

Avg Swap Used MB	The average amount of memory swapped, in megabytes.*
CPU Threads	The number of CPU threads.
Last Update	The date and time that the data in the table was last updated.

VMware Machines View

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](#)” 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](#)”, which shows current and historical utilization and performance metrics for a single virtual machine.

VMware Virtual Machines Table

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](#)” display.

VMware Virtual Machines Table 13-Jun-2019 11:28 No Alerts DATA

Server:

VMs: **15**

Virtual Machines

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 ▼

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.
- VMs:** The total number of virtual machines (rows) in the table.

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.
- Name** The name of the vSphere virtual machine.
- Alert Level** 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.

Heartbeat Status	<p>Displays whether or not the virtual machine has a heartbeat.*</p> <ul style="list-style-type: none"> ● 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.
Power State	Displays whether or not the virtual machine is powered on.*
Guest State	The state of the guest operating system.*
Number CPUs	The number of CPUs used by the virtual machine.*
VM Used CPU %	The percentage (%) of CPUs used.*
VM Memory Used (%)	The percentage (%) of memory used by the virtual machine.*
Mem Used	The amount of used memory, in megabytes.*
Memory Total MB	The total amount of memory, in megabytes.*
VM Disk Reads	The amount of data being read from the disk per second, in kilobytes.*
VM Disk Writes	The amount of data being written to the disk per second, in kilobytes.*
VM Net In KB/s	The amount of network data received per second, in kilobytes.*
VM Net OUT KB/s	The amount of network data transmitted per second, in kilobytes.*
Pkts Dropped IN %	The percentage of incoming packets that have been lost.*
Pkts Dropped OUT %	The percentage of outgoing packets that have been lost.*
Packets IN	The total number of incoming packets.*
Packets OUT	The total number of outgoing packets.*
Packet IN Dropped	The number of incoming packets that were dropped.*
Packets OUT Dropped	The number of outgoing packets that were dropped.*
Host Name	The name of the host.*
Guest Host Name	The name of the guest host.*
Guest IP Address	The IP address of the guest.*
Guest Operating System	The operating system used by the guest.*

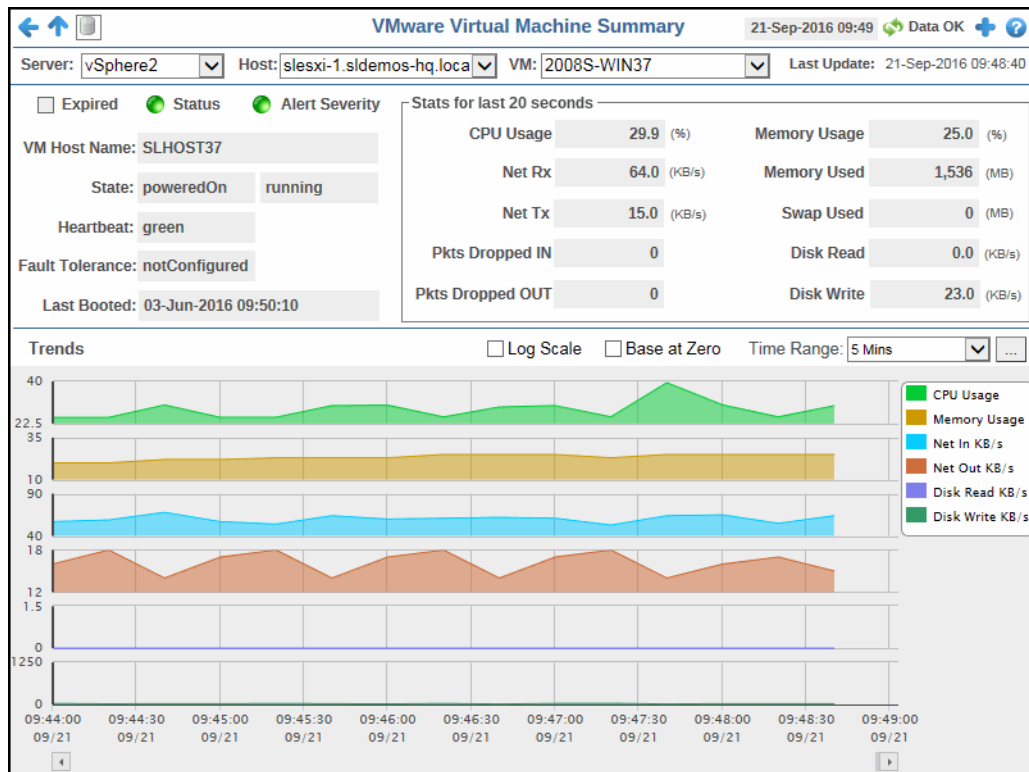
Connection State	The state of the current connection (connected/notConnected).*
Fault Tolerance	Displays whether or not fault tolerance is configured (configured/notConfigured).*
VM Tools Running Status	Displays whether or not the guest's tools are running (guestToolsRunning/guestToolsNotRunning).*
VM Tools Version Status	Displays the version status of the VMware tools installed on the guest operating system.*
Last Booted	The date and time the virtual machine was last rebooted.*
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.
Time Stamp	The date and time the row data was last updated.

VMware Virtual Machine Summary

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](#)" 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.



Filter By:

- Server** Select the server containing the virtual machine for which you want to view data.
- VM** Select the virtual machine for which you want to view data.

Fields and Data:

- Top Disk Utilization %** The maximum percentage of disk utilization across all disks used by the virtual machine.
- CPU %** The percentage of average CPU used on the virtual machine.*
- Memory Utilization %** The percentage of memory utilization on this machine.
- Network Rx/Tx KB/s** The rate of network received and transmitted data, in kilobytes per second, by this virtual machine.
- Dropped Packets Rx/Tx** The total number of dropped packets that were received and transmitted on this virtual machine.
- Disk Read/Write KB/s** The rate of disk reads and writes, in kilobytes per second, on this virtual machine.

Trend Graphs**Utilization**

CPU % -- Traces the percentage of CPU used.

Memory % -- Traces the amount of memory used.

Disk

Disk Reads KB/s -- Traces the amount of data being read from the disk, in kilobytes per second.

Disk Writes KB/s -- Traces the amount of data being written to the disk, in kilobytes per second.

CPU % -- Traces the percentage of CPU used.

Memory % -- Traces the percentage of memory used.

Network

Network In KB/s: The amount of network data received per second, in kilobytes per second.

Network Out KB/s: The amount of network data transmitted per second, in kilobytes per second.

CPU % -- Traces the percentage of CPU used.

Memory % -- Traces the percentage of memory used.


Network Loss

Dropped In Packets -- Traces the number of incoming packets that were dropped.

Dropped Out Packets -- Traces the number of outgoing packets that were dropped.

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 Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

Heartbeat Status	Displays whether or not the virtual machine has a heartbeat.* <ul style="list-style-type: none"> ● Red indicates that heartbeating has stopped. ● Yellow indicates that there might be a problem. ● Grey indicates that heartbeat status is disabled. ● Green indicates that heartbeat status is OK.
Pkts Dropped IN %	The percentage of incoming packets that were dropped.*
Critical/Warning	The number of critical and warning alerts.
Pkts Dropped OUT %	The percentage of outgoing packets that were dropped.*
Guest State	Displays whether or not the virtual machine is running.*
Top Disk Free %	The maximum percentage of free disk space across all disks used by the virtual machine.
Last Update	The date and time that the data in the table was last updated.

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.

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:

Note: This document assumes familiarity with the products monitored. For additional details, refer to vendor documentation.

The following Kafka Views and their associated displays are available in the Monitor. This section describes the Monitor displays and includes:

- ["Kafka Overview"](#): Describes the Kafka Overview display.
- ["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 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 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 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 Overview

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

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" 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" display, which allows you to view performance metrics for all servers on a particular cluster.

Kafka Clusters Table

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"](#) 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	▲	61	3	3	1	0	

Fields and Data:

Clusters

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:

Cluster

The name of the cluster.*

Alert Level

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.

Alert Count

The total number of alerts for the host.

Brokers Monitored

The current number of brokers being monitored for the cluster.*

Brokers Running

The number of brokers currently running on the cluster.*

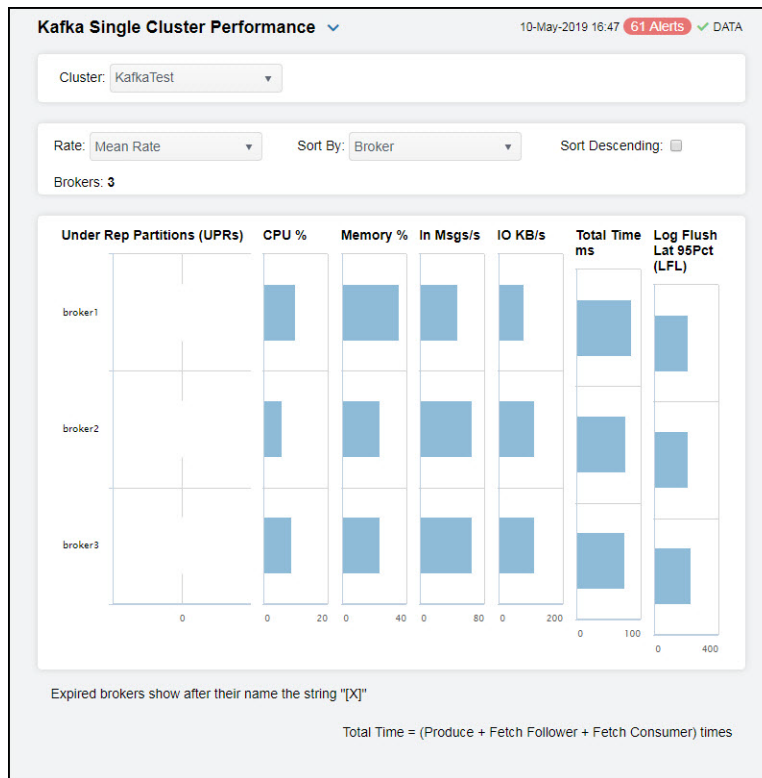
Active Controllers

The number of active controllers on the cluster.*

Offline Partitions	The number of partitions without an active leader on the cluster.*
Under Replicated Partitions	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.*
% Max Deviation in Partitions	The percentage of maximum deviation in partition count.*
% Max Deviation in Leaders	The percentage of maximum deviation in leader count.*
Zookeepers	The number of zookeepers on the cluster.*
Alive Connections	The number of connections on the zookeepers in the cluster.*
Outstanding Requests	The number of outstanding requests on the zookeepers in the cluster.
Received Packets/s	The rate of packets received on the zookeepers in the cluster.*
Sent Packets/s	The rate of packets sent by the zookeepers in the cluster.*
Consumed Bytes/s	The rate of bytes being consumed by the consumers.
Consumed Records/s	The rate of records being consumed by the consumers.
Consumers	The number of consumers on the cluster.*
Producer In Bytes/s	The rate of incoming bytes for the producers.*
Producer Out Bytes/s	The rate of outgoing bytes for the producers.*
Producer Sent Records/s	The rate of records being sent for the producers.*
Producers	The number of producers on the cluster.

Kafka Single Cluster Performance

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.

**Filter By:**

The display might include these filtering options:

- | | |
|------------------------|--|
| Cluster | Select the cluster for which you want to show data in the display. |
| Rate | Specify the rate for which you want to show data in the display. |
| 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 (servers) will be sorted in ascending or descending order using the option you select from this drop down. For example, if you select InMsgs/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 InMsgs/s will be listed at the top followed by the server with the next most InMsgs/s , and so on. |
| Sort Descending | When toggled on, the servers listed in the display are sorted in descending order based on the selected metric in the Sort By drop down. When toggled off, the servers are listed in ascending order. |

Fields and Data:

- | | |
|------------------------------------|--|
| Brokers | The number of brokers contained in the selected cluster. |
| Under Rep Partitions (UPRs) | Lists the number of partitions that are under-replicated or offline on each broker in the cluster. |
| CPU % | Lists the percentage of CPU used by the broker. |
| Memory % | Lists the percentage of memory used by the broker. |

In Msgs/s	Lists the rate of incoming messages (per second) for each broker in the cluster.
IO KB/s	Lists the rate of incoming kilobytes (per second) for each broker in the cluster.
Total Time ms	Lists the total time (in milliseconds) taken to service a request.*
Log Flush Latency 95Pct (LFL)	Lists the 95th percentile value for the log flush latency for each broker on the cluster.

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. Clicking **Kafka Brokers** from the left/navigation menu opens the “[Kafka Brokers Table](#)” 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](#)” display, which allows you to view performance metrics for all servers on a particular cluster.
- **Single Broker Summary:** Opens the “[Kafka Single Broker Summary](#)” 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](#)” display, which contains JVM runtime data for a single broker.
- **Single Broker Topics Summary:** Opens the “[Kafka Single Broker Topics Summary](#)” display, which contains topic data for a single broker.
- **Single Broker Topics Lag Summary:** Opens the “[Kafka Single Broker Topic Lag Summary](#)” display, which displays the lag per topic in a bar graph format and lists the lag per topic for the broker.

Kafka Brokers Table

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](#)” display and view metrics for that particular broker. Toggle between the commonly accessed displays by clicking the drop down list on the display title.

Kafka Brokers Table 15-May-2019 10:40 No Alerts ✓ DATA

Cluster: KafkaTest

Brokers: 3

Cluster	Broker	Broker ID	Alert Level	Alert Count	Broker State
KafkaTest	broker1	0	✓		RunningAsBroker
KafkaTest	broker2	1	✓		RunningAsBroker
KafkaTest	broker3	2	✓		RunningAsBroker

Filter By

Cluster Select the cluster for which you want to view data.

Brokers 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:

Cluster The name of the cluster.

Broker The name of the broker.

Broker ID The broker ID for the server.

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.

Broker State The current state of the kafka broker.*

Active Controller Denotes whether the broker is an active controller.*

Leader Count The number of leaders on the broker.*

Partitions The number of partitions on the broker.*

Offline Partitions The number of partitions without an active leader on the broker.*

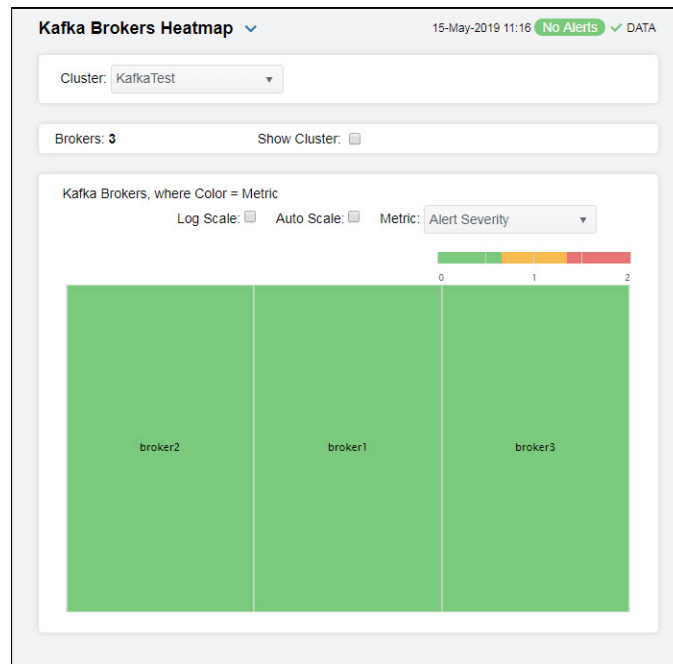
Under Replicated Partitions	The number of partition replicas that are out of sync (total number of replicas minus the total number of in-sync replicas) on the broker.*
Preferred Replica Imbalance Count	The number of topics whose replicas are not balanced on the broker.*
Purgatory Size Fetch	The number of fetch requests currently in purgatory (and waiting to be satisfied).*
Purgatory Size Heartbeat	The number of requests in purgatory due to failed heartbeat tests.*
Purgatory Size Produce	The number of produce requests currently in purgatory (and waiting to be satisfied).*
Purgatory Size Rebalance	The number of changes that need to be propagated to the replicas so that the partitions are no longer in purgatory.*
Purgatory Size Topic	The number of requests (based on topics) currently in purgatory.*
Network Processor Avg % Idle	The average fraction of time the network processors are idle.*
Kafka Version	The current version of Kafka.*
JMX Connection String	The JMX connection string used.*
Connected?	Denotes whether or not the broker is connected.*
Expired	<p>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.</p> <p>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.</p>
Timestamp	The date and time the row data was last updated.

Kafka Brokers Heatmap

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"](#) display.



Filter:

Cluster Select the cluster for which you want to view data.

Fields and Data:

Brokers Displays the number of brokers found based on the filter and that are displayed in the heatmap.

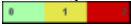




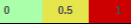
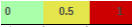
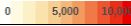
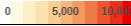
Show Cluster Select this check box to display the names of the cluster at the top of each rectangle in the heatmap.

Heatmap

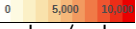
Log Scale Select this check box to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically 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.

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 Scale** is not selected.

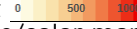
Metric	Select the metric driving the heatmap display. The default is Alert Severity . Each Metric has a color gradient bar that maps values to colors. The heatmap organizes the servers by host, where each rectangle represents a server. Mouse-over any rectangle to display the current values of the metrics for the broker. Click on a rectangle to drill-down to the associated "Kafka Single Broker Summary" display for a detailed view of metrics for that particular broker.
Alert Severity	<p>The current alert severity. Values range from 0 - 2, as indicated in the color gradient  bar, where 2 is the highest Alert Severity:</p> <ul style="list-style-type: none">  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	<p>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.</p>
Under Replicated Partitions	<p>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.</p> <p>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.</p>
Offline Partitions	<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 KafkaBrokerOfflinePartitionCnt. The middle value in the gradient bar indicates the middle value of the range.</p> <p>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.</p>
Msgs In Per Sec	<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 KafkaBrokerMsgsInPerSec. The middle value in the gradient bar indicates the middle value of the range.</p> <p>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.</p>
Bytes In Per Sec	<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 KafkaBrokerBytesInPerSec. The middle value in the gradient bar indicates the middle value of the range.</p> <p>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.</p>

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 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.

Log Flush Latency 95 Pct

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.

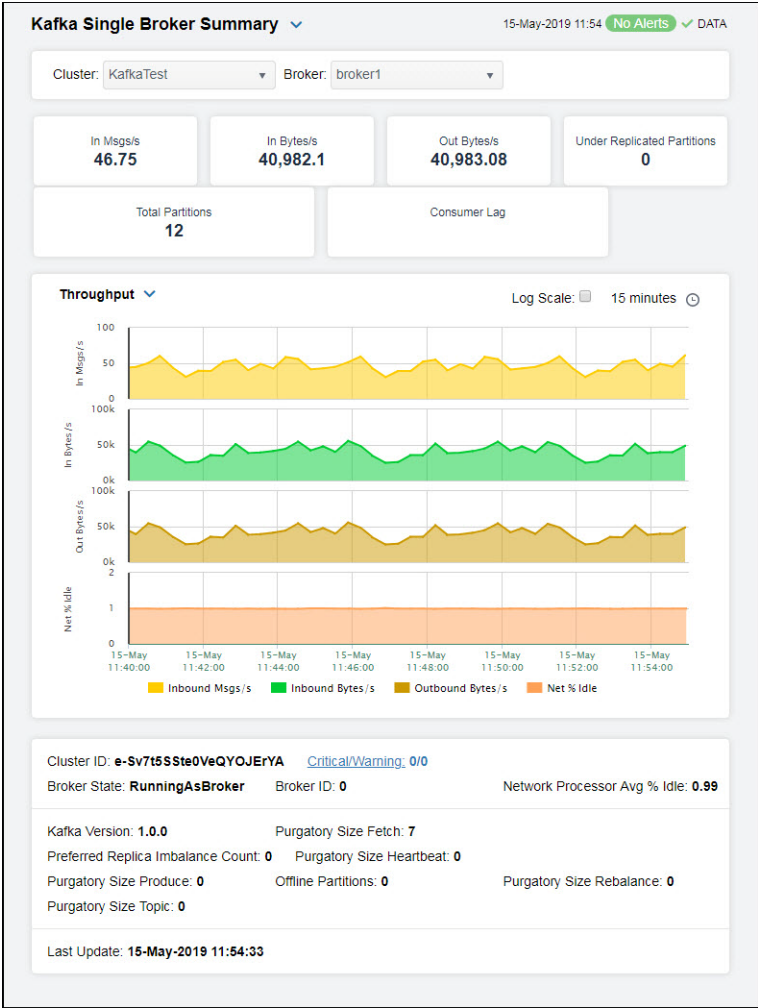
When **Auto Scale** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

Kafka Single Broker Summary

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](#)” 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.



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.

Fields and Data

- In Msgs/s** The number of incoming messages per second.
- In Bytes/s** The number of incoming bytes per second.
- Out Bytes/s** The number of outgoing bytes per second.
- Under Replicated Partitions** The number of partition replicas out of sync on the broker.
- Total Partitions** The total number of partitions on the broker.
- Consumer Lag** The aggregated consumer lag for all topics on the broker.

Trend Graphs

Throughput

- Inbound Msgs/s** -- traces the number of incoming messages per second.
- Inbound Bytes/s** -- traces the number of incoming bytes per second.
- Outbound Bytes/s** -- traces the number of outgoing bytes per second.
- Net % Idle** -- traces the average fraction of time the network processors are idle.


Partitions

- Offline Partitions** -- traces the number of offline partitions.
- Under Replicated Partitions** -- traces the number of partition replicas out of sync on the broker.
- Active Controllers** -- traces whether or not the broker is/was an active controller.

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 Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

Cluster ID	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.
Broker State	The current state of the Kafka broker.
Network Processor Avg % Idle	The average fraction of time the network processors are idle.*
Critical/Warning	The number of critical and warning alerts.
Broker ID	The broker ID for the server.
Kafka Version	The current version of Kafka installed on the broker.
Preferred Replica Imbalance Count	The number of topics whose replicas are not balanced on the broker.*
Purgatory Size Produce	The number of produce requests currently in purgatory (and waiting to be satisfied).*
Purgatory Size Rebalance	The frequency with which the partition rebalance check is triggered by the controller.*
Purgatory Size Fetch	The number of fetch requests currently in purgatory (and waiting to be satisfied).*

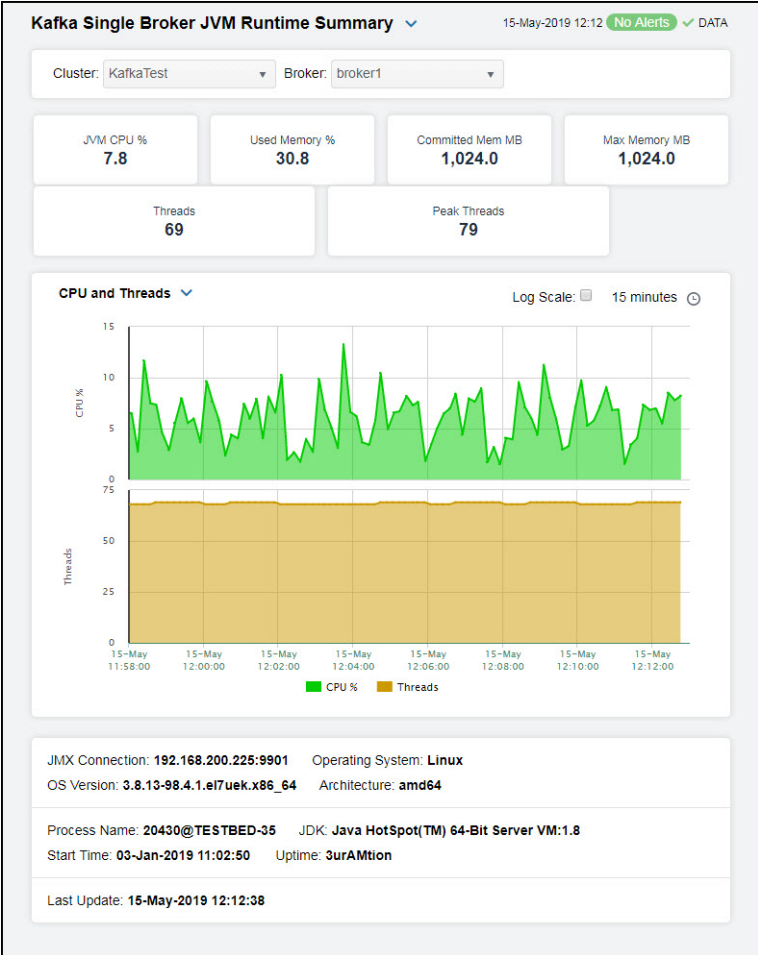
Purgatory Size Heartbeat	The number of requests in purgatory due to failed heartbeat tests.*
Offline Partitions	The number of partitions on the broker that are currently offline.*
Purgatory Size Topic	The number of requests (based on topics) currently in purgatory.*
Last Update	The date and time of the last data update.

Kafka Single Broker JVM Runtime Summary

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"** 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.



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.

Fields and Data

- JVM CPU %** The percentage of CPU used of this broker as JVM.
- Used Memory %** The percentage of memory used of this broker as JVM.
- Committed Mem MB** The committed heap memory, in megabytes, of this broker as JVM.
- Max Memory MB** The maximum heap memory, in megabytes, of this broker as JVM.
- Threads** The number of threads running in the broker.
- Peak Threads** The peak number of threads running in the broker.

Trend Graphs**CPU and Threads**

CPU % -- traces the percentage of CPU used of this broker as JVM.

Threads -- traces the number of threads running in the broker.

Heap Memory

Max Mem MB-- traces the maximum heap memory, in megabytes, of this broker as JVM.


Committed Mem MB -- traces the committed heap memory, in megabytes, of this broker as JVM.

Used Mem MB -- traces the memory used by 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.

Time Settings Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

JMX Connection	The name of the JMX connection.*
Architecture	The type of processor being used.*
Operating System	The operating system installed on the broker.*
OS Version	The version number of the operating system.*
Process Name	The name of the process.*
Start Time	The date and time when the broker was started.*
JDK	The JDK version number.*
Uptime	The amount of time the broker has been up and running.*
Last Update	The date and time of the last data update.

Kafka Single Broker Topics Summary

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”](#) display and view metrics for that particular topic. Toggle between the commonly accessed displays by clicking the drop down list on the display title.

Topic	In Bytes/s	Out Bytes/s	Rejected Bytes/s	Failed Fetched Requests/s	Failed Pr 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	

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.
- Rate** 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.

Topics The total number of topics listed in the table.

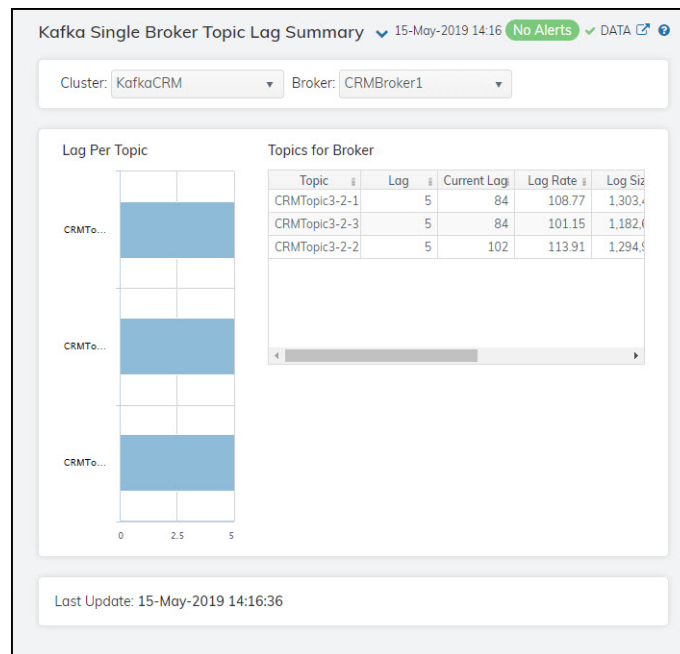
Metrics by Topic for Selected Broker Table

Topic	Lists the name of the topic.
In Bytes/s	The rate of incoming bytes
Out Bytes/s	The rate of outgoing bytes.
Rejected Bytes/s	The rate of rejected bytes.
Failed Fetch Requests/s	The rate of failed fetch requests.
Failed Produce Requests Per Sec	The rate of failed produce requests.
Fetches Msg Conversions/s	The rate of fetched message conversions, per second.
In Msgs/s	The rate of incoming messages
Produced Msgs Conversions/s	The rate of produced message conversions, per second.
Total Fetch Requests/s	The rate of total fetch requests.
Total Produce Requests/s	The rate of total produce requests.

Kafka Single Broker Topic Lag Summary

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](#)” 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.



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.

Lag Per Topic Bar Graph Displays the lag per topic in a bar graph format.

Topics for Broker Table

topic	The name of the topic.
Lag	The difference between the current consumer position in the partition and the end of the log.*
Current Lag	The difference in the amount of lag from the previous polling period to the current polling period.*
Lag Rate	The rate of change in the amount of lag.*
Log Size	The current number of messages in the log.*
Partitions	The number of partitions containing the topic.
Time Stamp	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. Clicking **Kafka Zookeepers** from the left/navigation menu opens the ["Kafka Zookeepers Table"](#) 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"](#) display, which allows you to view performance metrics for all servers on a particular cluster.
- **Single Zookeeper Summary**: Opens the ["Kafka Single Zookeeper Summary"](#) 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"](#) display, which contains JVM Runtime metrics for a single zookeeper.

Kafka 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. 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"](#) 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	zookeeper1		✓		

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.

Zookeepers The number of zookeepers that were found in the selected cluster.

Table

Cluster The name of the cluster.

Zookeeper The name of the zookeeper.

URL The URL associated with the zookeeper.

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.
Role	The role of the zookeeper (Leader or Follower).*
Max Request Latency ms	The longest amount of time taken to respond to a client request (in milliseconds) on the zookeeper since the last polling update.*
Avg Request Latency ms	The average amount of time taken to respond to a client request (in milliseconds) on the zookeeper since the last polling update.*
Min Request Latency	The least amount of time taken to respond to a client request (in milliseconds) on the zookeeper since the last polling update.*
Alive Connections	The number of clients connected to the zookeeper.*
Outstanding Requests	The number of queued requests.*
Node Count	The total number of nodes.*
Watch Count	The number of watchers set up over the zookeeper nodes.*
Packets Recvd	The number of packets received.*
Packets Sent	The number of packets sent.*
Current Packets Recvd	The increase in the amount of packets received by the zookeeper (from the previous polling period to the current polling period).*
Current Packets Sent	The increase in the amount of packets sent from the zookeeper (from the previous polling period to the current polling period).*
Received Packets/s	The rate at which packets are being received by the zookeeper.*
Sent Packets/s	The rate at which packets are being sent by the zookeeper.*
Max Client Cnxns Per Host	The maximum number of connections allowed from each host.*
Max Session Timeout	The maximum allowed session timeout allowed for registered consumers.*
Min Session Timeout	The minimum allowed session timeout allowed for registered consumers.*
Kafka Version	The current version of Kafka being used.*
Client Port	Lists the client's port.*
JMX Connection String	Lists the connection string.*
Connected	Denotes whether or not the zookeeper is connected.

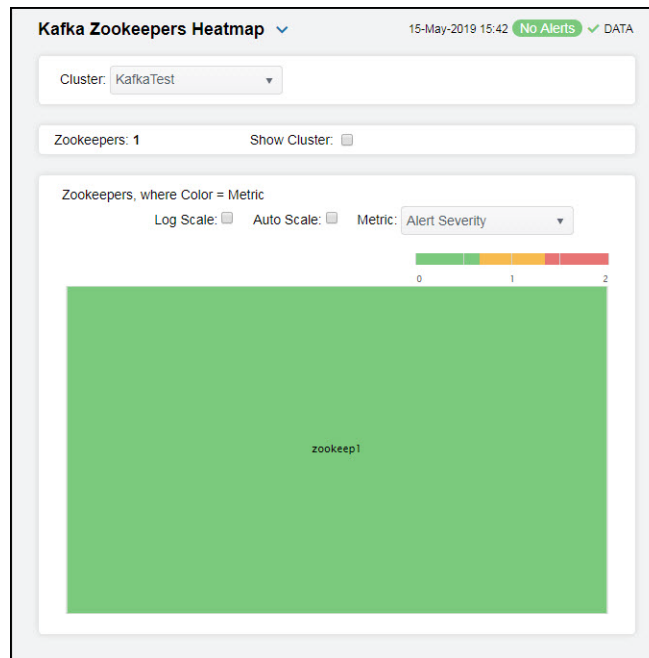
Expired	<p>When checked, performance data in the row has not been received within the time specified (in seconds) in the Expire Time field in the RTView Configuration Application > (KAFKAMON-LOCAL/Project Name) > Solution Package Configuration > Apache Kafka > DATA STORAGE > Duration > Expire Time property. 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.</p> <p>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.</p>
Start Time	The date and time the zookeeper was started. *
Source	The source of the zookeeper.
Timestamp	The date and time the row data was last updated.

Kafka Zookeepers Heatmap

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"](#) display.

**Filter By:****Cluster**

Select the cluster for which you want to view data.

Zookeepers

The number of zookeepers found on the cluster and listed in the heatmap.

Show Cluster

Select this check box to display the name of the cluster at the top of each rectangle in the heatmap.

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.






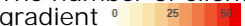

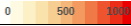
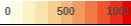
Auto Scale

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 Scale** is not selected.

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 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"](#) display for a detailed view of metrics for that particular zookeeper.

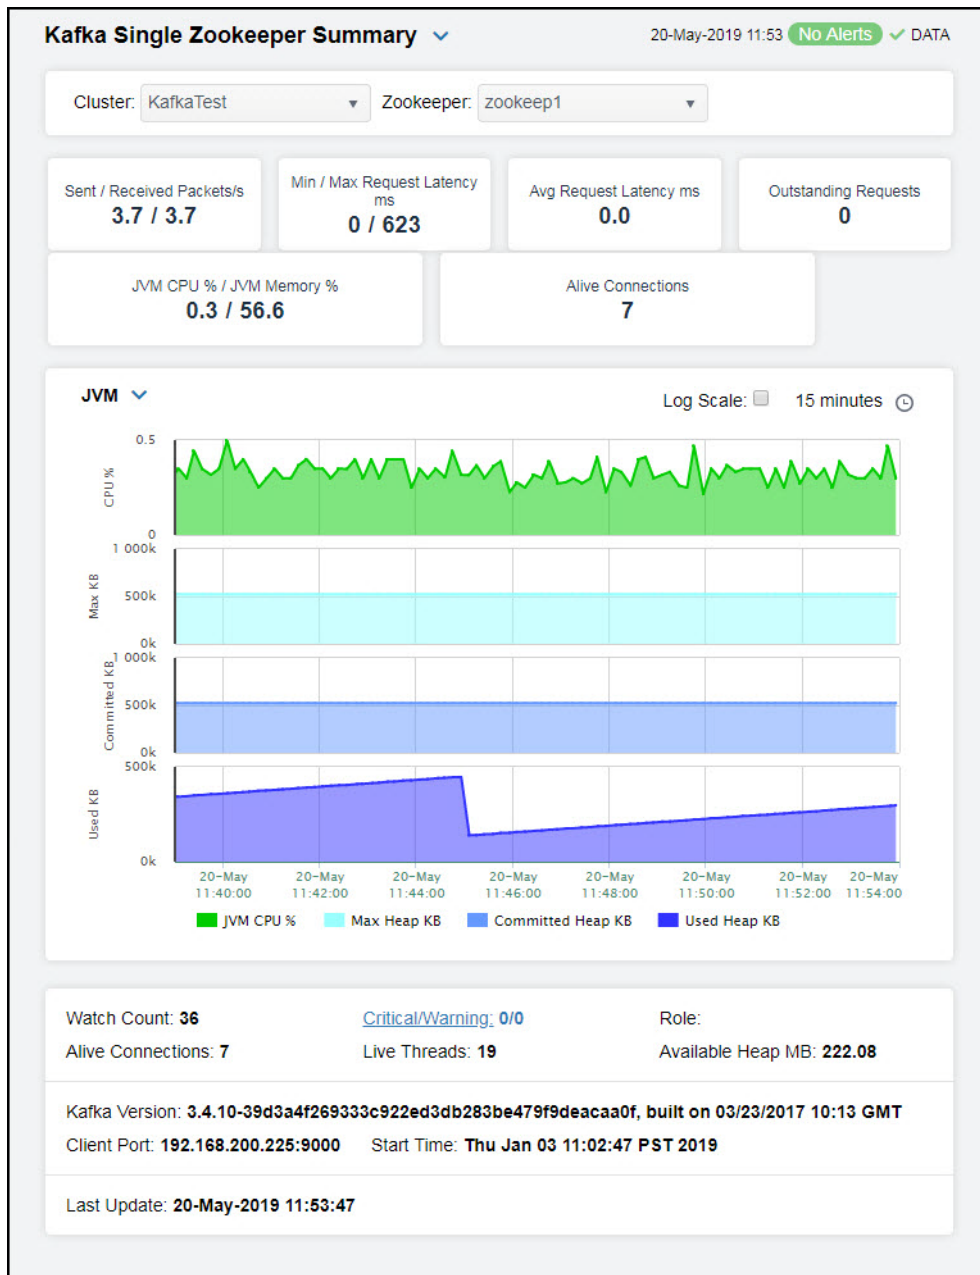
Alert Severity	<p>The current alert severity. Values range from 0 - 2, as indicated in the color gradient  bar, where 2 is the highest Alert Severity:</p> <ul style="list-style-type: none">  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	<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 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.</p>
# Alive Connections	<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 KafkaZookeeperNumAliveConns. The middle value in the gradient bar indicates the middle value of the range.</p> <p>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.</p>
Outstanding Reqs	<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 KafkaZookeeperOutstandingReqs. The middle value in the gradient bar indicates the middle value of the range.</p> <p>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.</p>
Packets In Per Sec	<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 KafkaZookeeperRatePktsRcvd. The middle value in the gradient bar indicates the middle value of the range.</p> <p>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.</p>
Packets Out Per Sec	<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 KafkaZookeeperRatePktsSent. The middle value in the gradient bar indicates the middle value of the range.</p> <p>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.</p>

Kafka Single Zookeeper Summary

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"](#) 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.



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 in the display.

Zookeeper Select the zookeeper for which you want to view data in the display.

Fields and Data

Sent/ Received Packets/s	The rate at which packets are being sent and received (per second) by the zookeeper.*
Min/Max Request Latency ms	The shortest and longest amount of time taken to respond to a client request (in milliseconds) on the zookeeper since the last polling update.*
Avg Request Latency ms	The average amount of time taken to respond to a client request (in milliseconds) on the zookeeper since the last polling update.*
Outstanding Requests	The number of queued requests.*
JVM CPU %/ JVM Memory %	The percentage of CPU being used by the JVM, and the percentage of heap memory used by the zookeeper.*
Alive Connections	The number of clients connected to the zookeeper.*

Trend Graphs

Latency

Max Request Latency ms -- traces the longest amount of time taken to respond to a client request, in milliseconds.

Avg Request Latency ms -- traces the average amount of time taken to respond to a client request, in milliseconds.

Min Request Latency ms -- traces the least amount of time taken to respond to a client request, in milliseconds.

Performance

Outstanding Requests -- traces the number of queued requests.

Sent Packets/s -- traces the rate at which packets are being sent (per second) by the zookeeper.

Rcvd Packets/s -- traces the rate at which packets are being received (per second) by the zookeeper.

Alive Connections -- traces the number of alive connections.

JVM

JVM CPU % -- traces the percentage of CPU being used by the JVM.

Max Heap KB -- traces the maximum amount of available heap, in kilobytes.

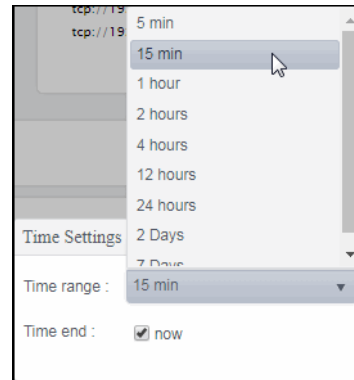
Committed Heap KB -- traces the amount of committed heap, in kilobytes.

Used Heap KB -- traces the highest amount of heap used, 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.

Time Settings Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

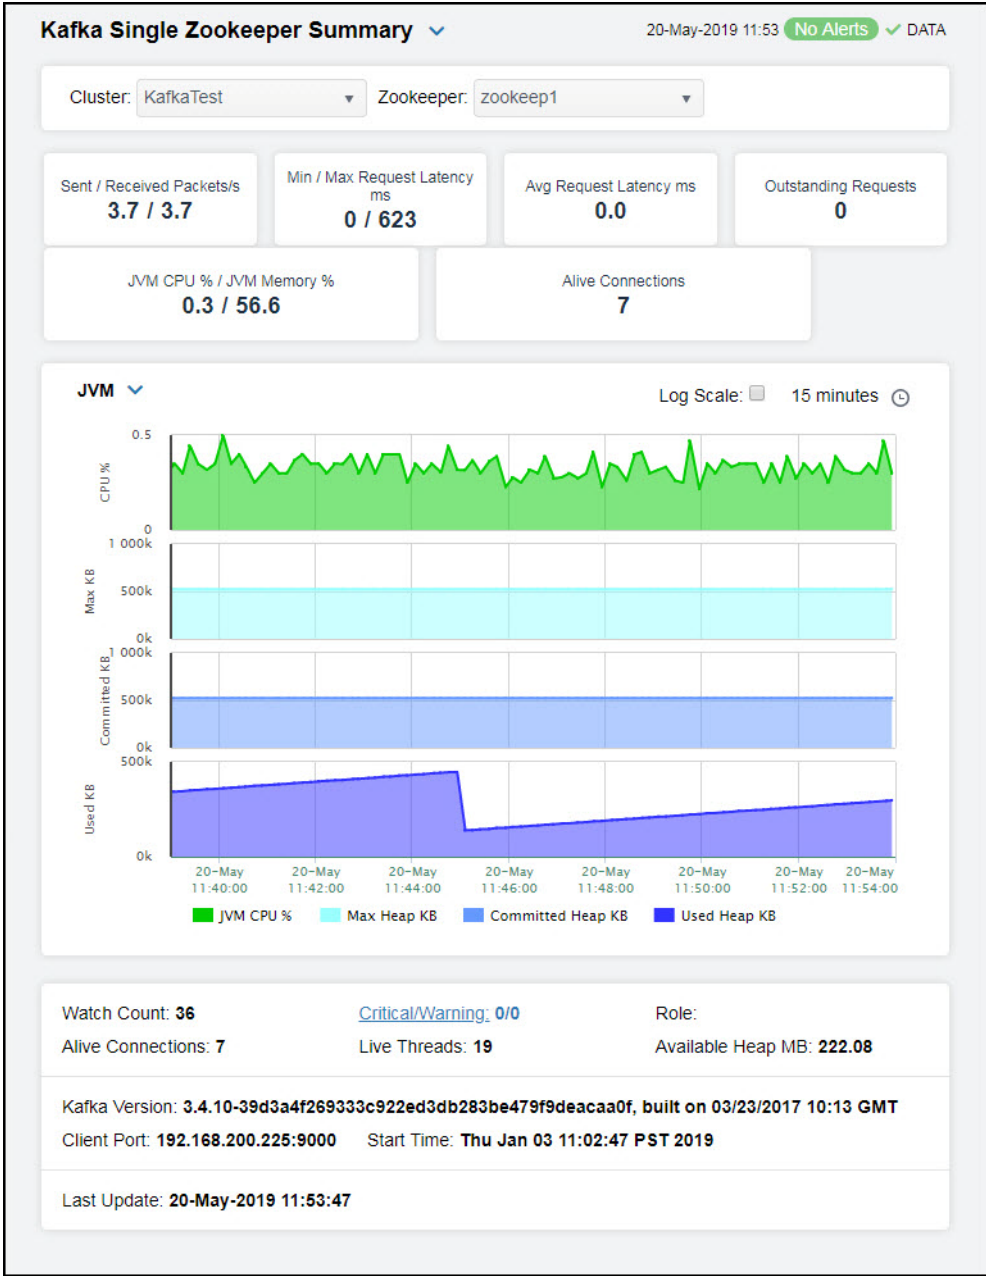
Watch Count	The number of watchers set up over the zookeeper nodes.*
Alive Connections	The number of clients connected to the zookeeper.*
Critical/Warning	The total number of critical and warning alerts.
Live Threads	The total number of live threads.*
Role	The zookeeper's role (Leader/Follower).*
Available Heap MB	The maximum amount of available heap, in megabytes.*
Kafka Version	The current version of Apache Kafka installed.*
Client Port	The client's IP address and port.*
Start Time	The date and time when the zookeeper was started.*
Last Update	The date and time of the last data update.

Kafka Single Zookeeper JVM Runtime Summary

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"](#) 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.



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.

Fields and Data

CPU % The percentage of CPU used by the zookeeper.

Used Memory % The percentage of memory used by the zookeeper as JVM.

Committed Mem MB The committed heap memory, in megabytes.

Max Memory MB The maximum amount of heap memory, in megabytes.

Threads The number of threads running on the zookeeper.

Peak Threads The peak number of threads running on the zookeeper.

Performance Trends Trend Graph

Traces the following:

CPU % -- traces the percentage of CPU being used by the JVM.


Max Mem MB -- traces the maximum amount of available heap.

Committed Mem MB -- traces the amount of committed heap memory.

Used Mem MB -- 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.

Time Settings Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

- JMX Connection** The name of the JMX connection. *
- OS Version** The version number of the operating systems. *
- Process Name** The name of the process. *
- Start Time** The date and time when the zookeeper was started. *
- Critical/Warning** The total number of critical and warning alerts.
- Architecture** The type of processor being used. *
- JDK** The date and time when the zookeeper was started. *
- Uptime** The amount of time the zookeeper has been up and running. *
- Operating System** The operating system installed on the zookeeper. *
- Last Update** The date and time of the last data update.

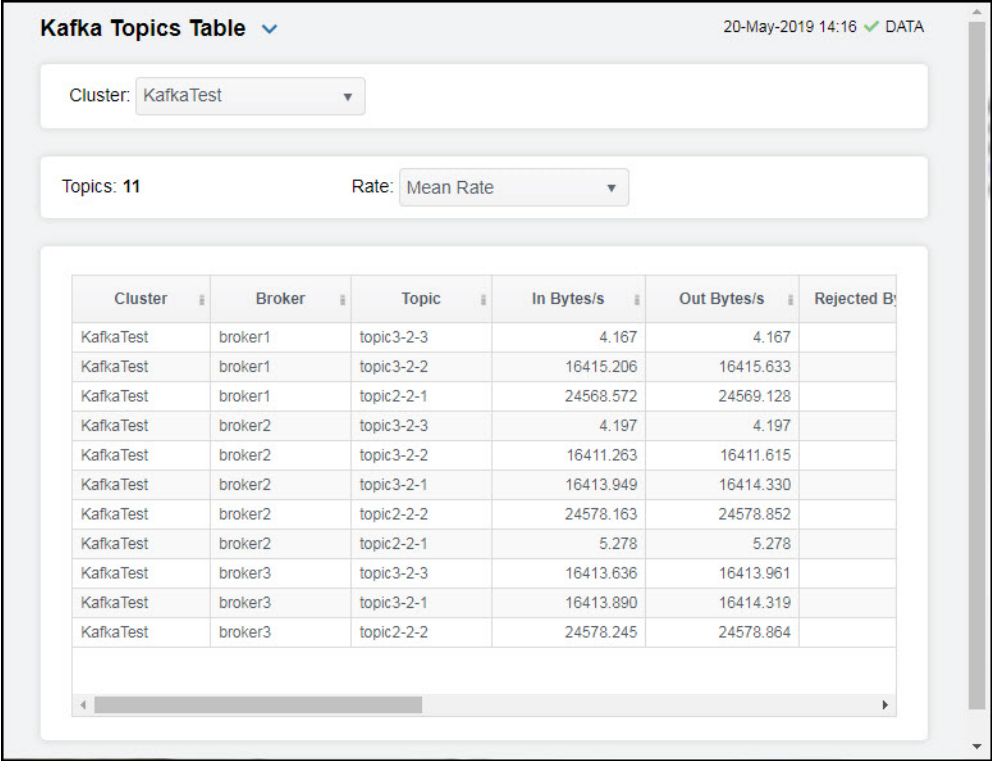
Kafka Topics View

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. Clicking **Kafka Topics** from the left/navigation menu opens the "[Kafka Topics Table](#)" 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](#)" display, which allows you to view performance metrics for all servers on a particular cluster.
- **Single Topic Summary**: Opens the "[Kafka Single Topic Summary](#)" 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](#)" 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](#)" 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](#)" display, which allows you to view performance metrics for all brokers for a particular topic.

Kafka Topics Table

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](#)" display and view metrics for that particular topic. Toggle between the commonly accessed displays by clicking the drop down list on the display title.



Filter By

Cluster Select the cluster for which you want to view data.

Topics The total number of topics listed in the table.

Rate 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.

Table:

Cluster Lists the name of the cluster.

Broker The name of the broker.

Topic Lists the name of the topic.

In Bytes/s The rate of incoming bytes.

Out Bytes/s The rate of outgoing bytes.

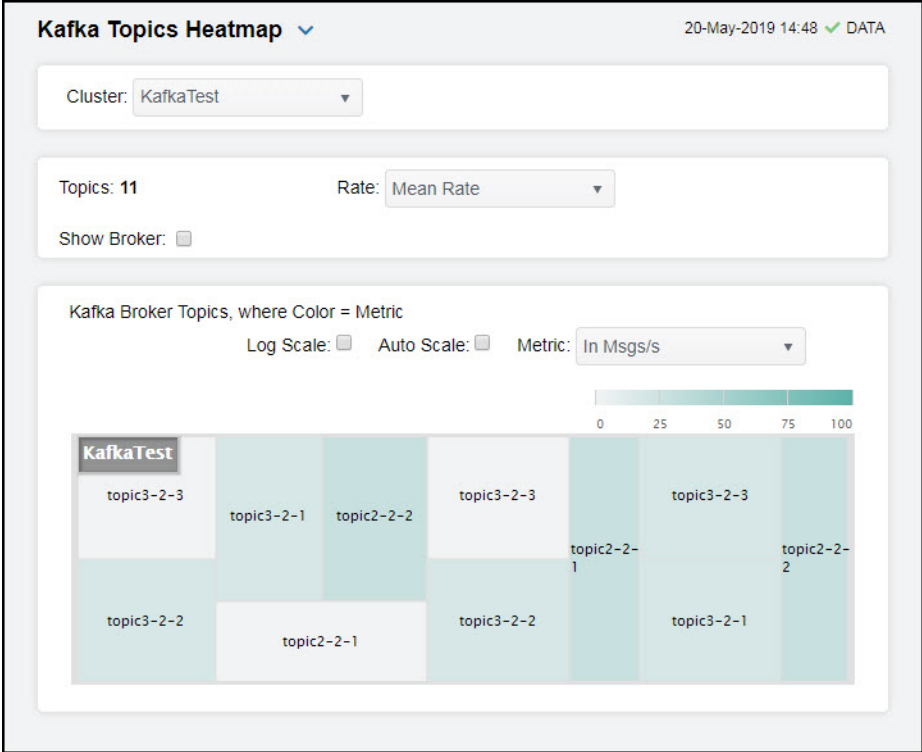
Rejected Bytes/s	The rate of rejected bytes.
Failed Fetched Requests/s	The rate of failed fetch requests.
Failed Produce Requests/s	The rate of failed produce requests.
Fetched Msg Conversions/s	The rate of fetched message conversions.
In Msgs/s	The rate of incoming messages.
Produced Msgs Conversions/s	The rate of produced message conversions.
Total Fetch Requests/s	The rate of total fetch requests.
Total Produce Requests/s	The rate of total produce requests.

Kafka Topics Heatmap

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"](#) display.



Fields and Data:

Cluster Select the cluster for which you want to view data.

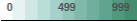

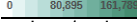
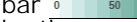
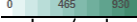
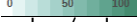
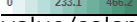
Count Lists the number of topics displayed in the heatmap.

Rate 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.

Show Broker Select this check box to display the names of the Brokers at the top of each rectangle in the heatmap.

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.

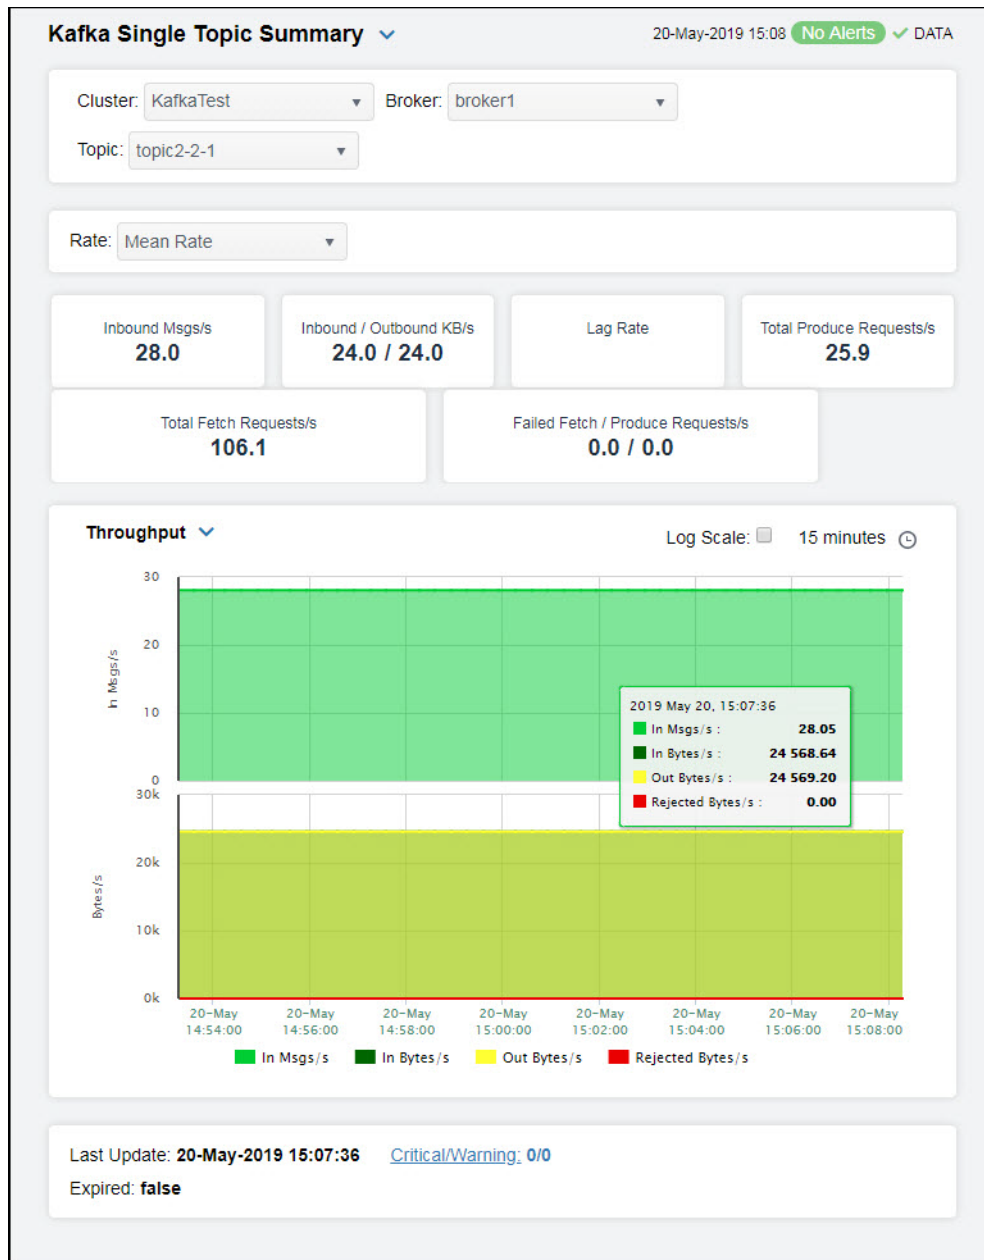
Auto Scale	Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value. Note: Some metrics auto-scale automatically, even when Auto Scale is not selected.
Metric	Select the metric driving the heatmap display. The default is In Msgs/s . 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" 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, 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 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

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](#)" 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.



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.
Rate	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.
Inbound Msgs/s	The rate of incoming messages, per second.
Inbound/ Outbound KB/s	The rate of inbound and outbound kilobytes, per second.
Lag Rate	The rate of change in the amount of lag.
Total Produce Requests/s	The rate of total produce requests.
Total Fetch Requests/s	The rate of total fetch requests.
Failed Fetch/ Produce Requests/s	The rate of failed fetch requests, and the rate of produce requests.
Trend Graphs	<p>Throughput</p> <ul style="list-style-type: none"> In Msgs/s -- traces the selected rate of incoming messages. In Bytes/s -- traces the selected rate of incoming bytes. Out Bytes/s -- traces the selected rate of outgoing bytes. Rejected Bytes/s -- traces the selected rate of rejected bytes. <p>Produce Requests</p> <ul style="list-style-type: none"> Total Produce Requests/s -- traces the selected rate of total produce requests. Failed Produce Requests/s -- traces the selected rate of failed produce requests. Lag Rate -- traces the rate of change in the amount of lag. <p>Fetch Requests</p> <ul style="list-style-type: none"> Total Fetch Requests/s-- traces the selected rate of total fetch requests. Failed Fetch Requests/s -- traces the selected rate of failed fetch requests. Lag Rate -- traces the rate of change in the amount of lag. <p>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.</p>

Time Settings Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

Last Update The date and time of the last data update.

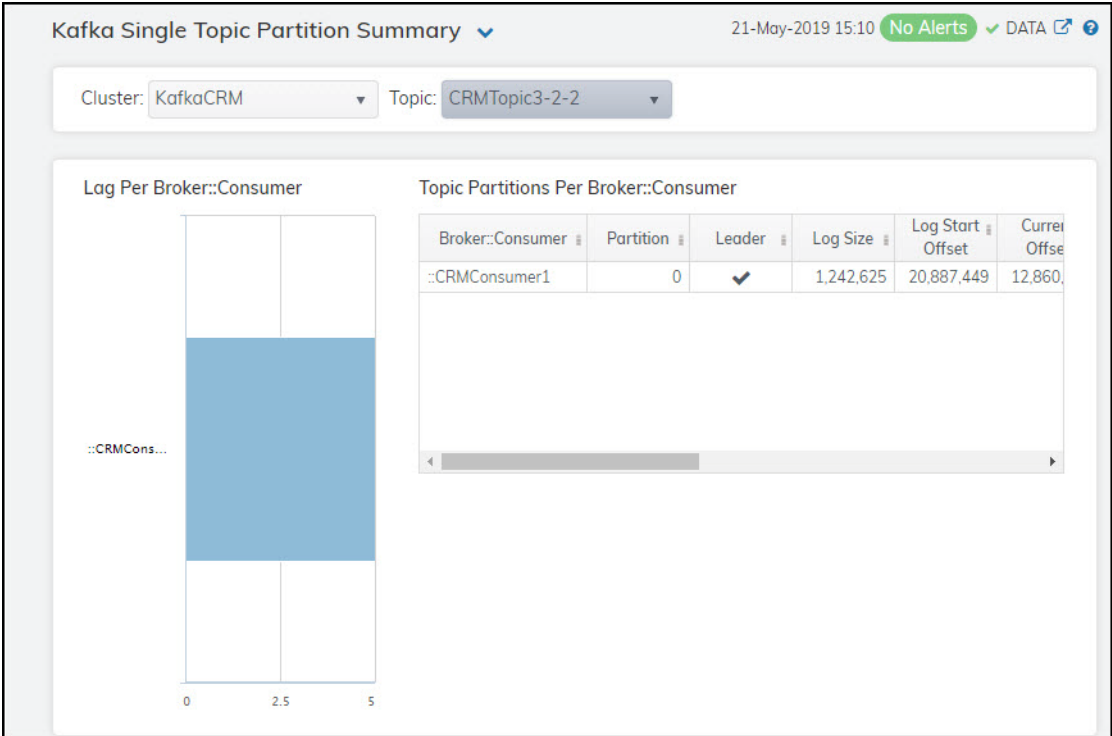
Critical/Warning The total number of critical and warning alerts.

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.

Kafka Single Topic Partition Summary

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.



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.
- Topic** Select the topic for which you want to see data.

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 Per Broker::Consumer Table

- Broker::Consumer** The name of the broker/consumer.
- Partition** The name of the partition.

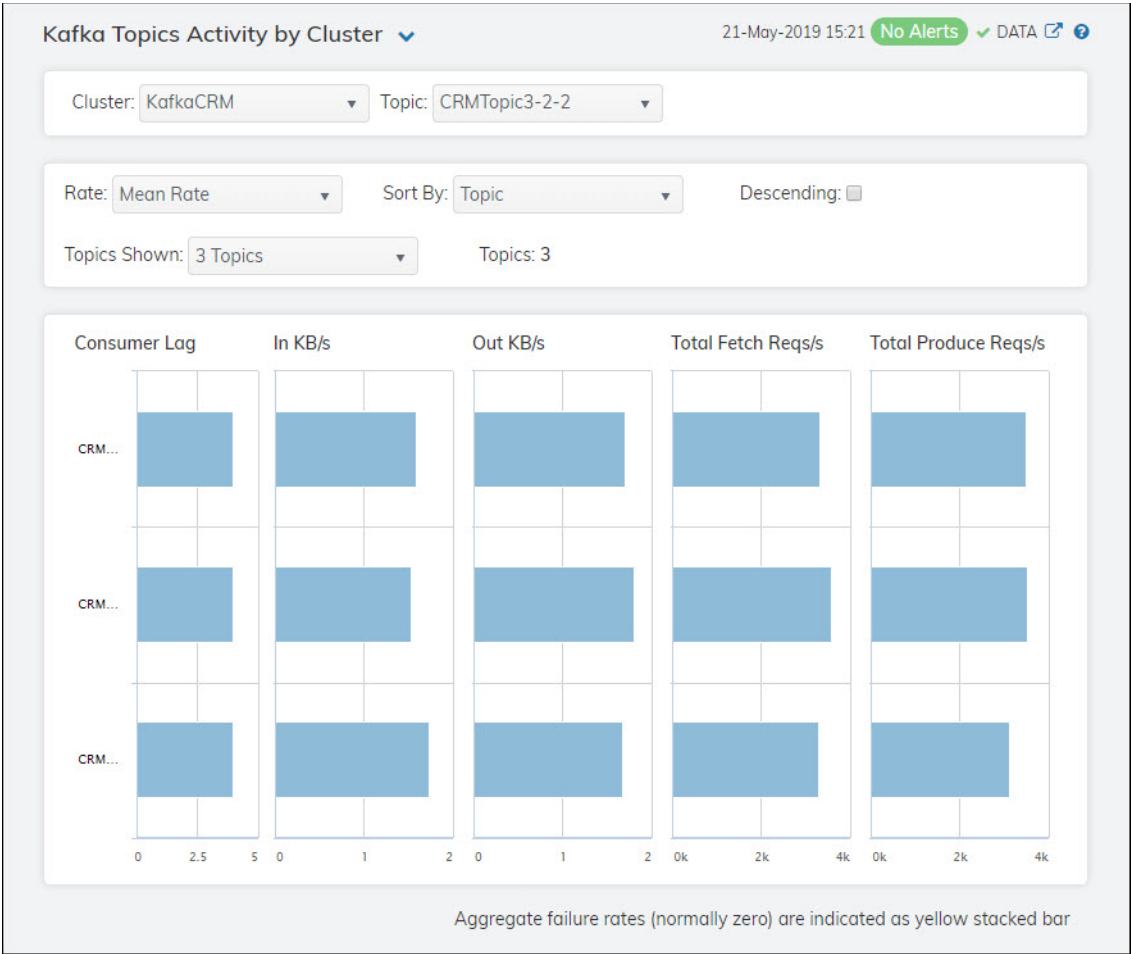
Leader	When checked, signifies that the broker is a leader on the partition.
Log Size	The current number of messages in the log.*
Log Start Offset	The offset of the first message written to the log.
Current Offset	The offset of the message currently being consumed.*
Log End Offset	The offset of the last message written to the log.*
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.*
Lag Rate	The rate of change in the amount of lag.*
Current Delta	The difference between the current consumer position in the partition from the previous polling period to the current polling period.*
Current Rate	The rate of change of the current consumer position.*
Log End Delta	The difference between the offset of the last message in the partition from the previous polling period to the current polling period.*
Log End Rate	The rate of change of the last message offset.*

Kafka Topics Activity by Cluster

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.



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.
- Topic** Select the topic for which you want to see data.

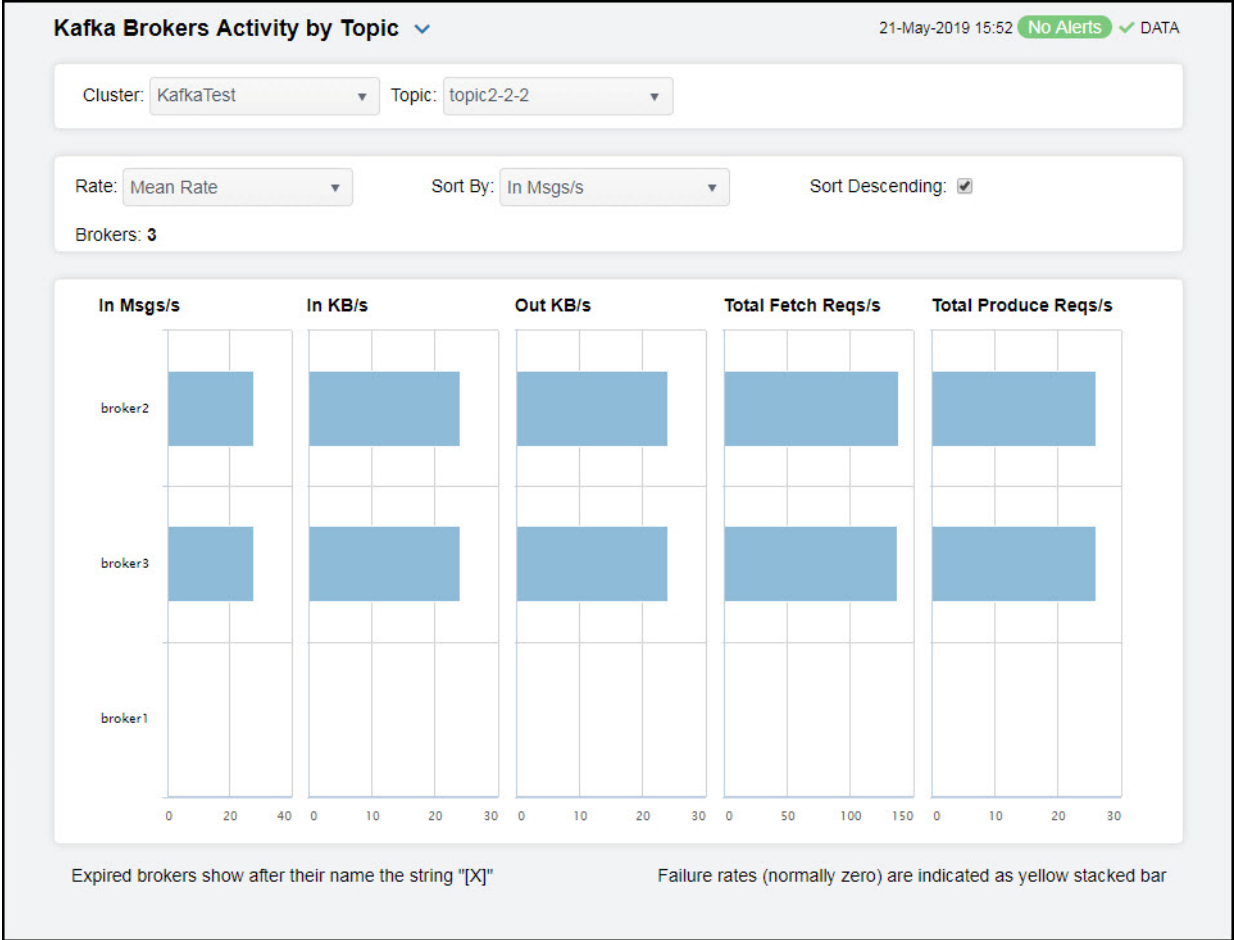
Fields and Data

Rate	<p>Select the option for which you want to view data:</p> <p>Mean Rate -- Select this option to view the average rate for each metric for the topics in the display.</p> <p>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.</p> <p>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.</p> <p>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.</p>
Sort By	<p>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 In KB/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 KB/s will be listed at the top followed by the topic with the next most In KB/s, and so on.</p>
Descending	<p>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.</p>
Topics Shown	<p>Select the number of topics that you want to view in the trend graphs below.</p>
Topics	<p>The number of topics found in the cluster.</p>
Graphs	
Consumer Lag	<p>The difference between the current consumer position in the partition and the end of the log.*</p>
In KB/s	<p>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.</p>
Out KB/s	<p>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.</p>
Total Fetch Reqs/s	<p>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.</p>
Total Produce Reqs/s	<p>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.</p>

Kafka Brokers Activity by Topic

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.



Filter By:

- Cluster**
- Topic**
- Rate**

Select the cluster for which you want to see data.

Select the topic for which you want to see data.

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 brokers 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 brokers 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 brokers 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 brokers in the display.

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 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.
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.
Brokers	The number of brokers found in the cluster with the associated topic.
Graphs	
In Msgs/s	The number of incoming messages per second. For example, if you select 1 Minute Avg from the drop down list, the average rate of incoming messages per second for 1 minute.
In KB/s	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.
Out KB/s	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/s	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/s	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.

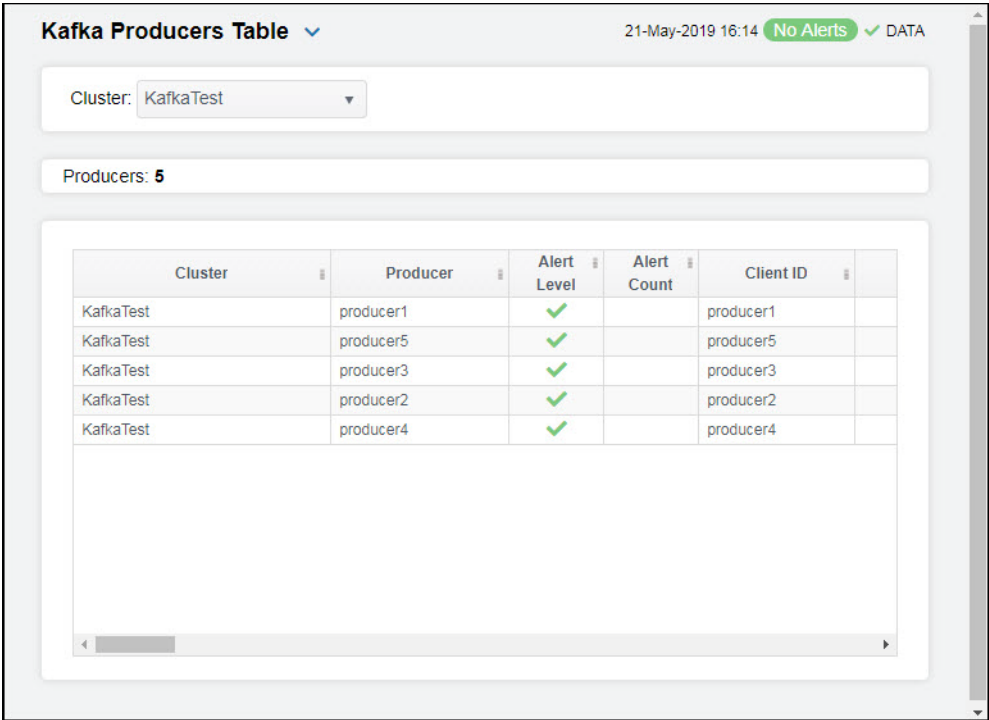
Kafka Producers View

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](#)” 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](#)” 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](#)” 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](#)” display, which contains current and historical JVM runtime metrics, as well as trend data, for a single producer.

Kafka 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. 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" display and view metrics for that particular producer. Toggle between the commonly accessed displays by clicking the drop down list on the display title.



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.

Producers The number of producers found on the selected cluster, and that are listed in the **All Producers Table**.

Table:

Cluster The name of the cluster.*

Producer The name of the producer.

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.
Client ID	The ID of the producer.*
Batch Size Average	The average batch size sent by the producer.*
Batch Size Max	The maximum number of messages that can be added to a batch before being sent to the event handler.*
Buffer Available Bytes	The number of available bytes in the buffer.*
Buffer Exhausted Rate	The average per-second number of record sends that are dropped due to buffer exhaustion.*
Buffer Total Bytes	The total number of bytes allowed in the buffer.*
Buffer Pool Wait Ratio	The fraction of time an appender waits for space allocation.*
Compression Avg Rate	The average compression rate of record batches.*
Connection Close Rate	The rate of connections being closed.*
Connections	The number of active connections.*
Connection Creation Rate	The rate of connections being created.*
Producer In Bytes/s	The average number of incoming bytes per second.*
IO Ratio	The rate of input/output operations.*
IO Time NS Avg	The average length of time the I/O thread spent waiting for a socket (in nanoseconds).*
IO Wait Ratio	The percent of time the producer was performing I/O operations while the CPU was idle.*
IO Wait Time ms Avg	The average length of time the I/O thread spent waiting for a socket (in milliseconds).*
Metadata Age	The age (in seconds) of the current producer metadata being used.*
Network IO Rate	The rate of input/output network operations.*
Producer Out Bytes/s	The average number of outgoing bytes per second.*
Produce Throttle Time Avg	The avg time (in milliseconds) a request was throttled by a broker.*

Produce Throttle Time Max	The maximum time (in milliseconds) a request was throttled by a broker.*
Record Errors/s	The average per-second number of record sends that resulted in errors for a topic.*
Avg Record Queue Time ms	The average time (in milliseconds) record batches spent in the record accumulator.*
Max Record Queue Time ms	The maximum time (in milliseconds) record batches spent in the record accumulator.*
Record Retry Rate	The average per-second number of retried record sends.
Producer Sent Records/s	The average number of records sent (per second) for a topic.*
Record Size Avg	The average record size.*
Record Size Max	The maximum record size.*
Records per Request Avg	The average number of records per request.*
Request Latency Avg	The average request latency (in milliseconds).*
Request Latency Max	The maximum request latency (in milliseconds).*
Request Rate	The average number of requests sent per second.*
Request Size Avg	The average request size.*
Request Size Max	The maximum request size.*
Requests In Flight	The current number of in-flight requests awaiting a response.*
Response Rate	The average number of responses received per second.*
Select Rate	The number of times the I/O layer checked for new I/O operations to perform per second.*
Waiting Threads	The number of user threads blocked waiting for buffer memory to enqueue their records.*
Jmx Connection	The JMX connection string.*
Kafka Version	The current version of Apache Kafka installed.*
Connected	Denotes whether or not the producer is connected.

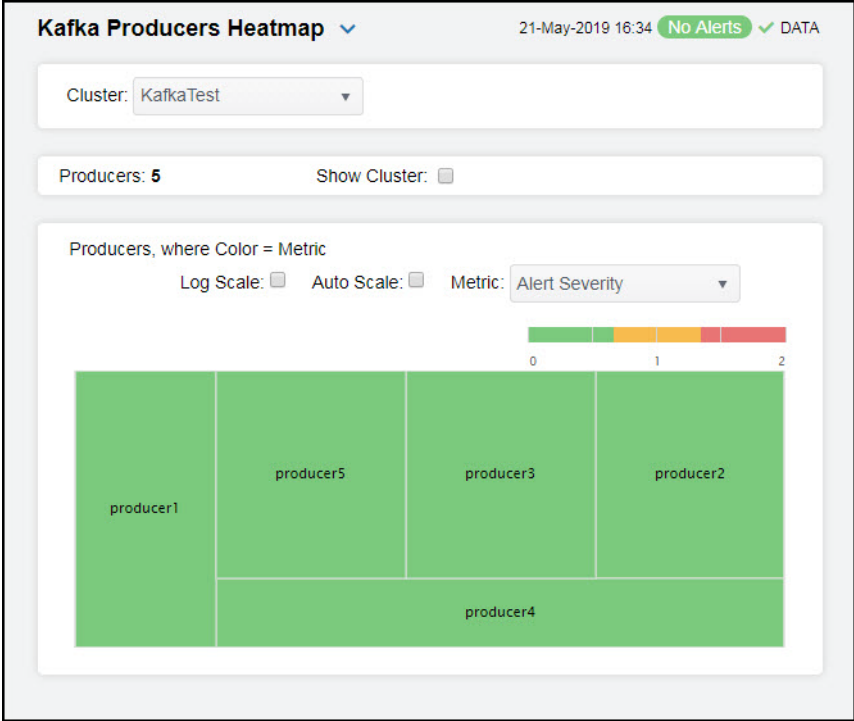
- Expired** When checked, performance data in the row has not been received within the time specified (in seconds) in the **Expire Time** field in the RTView Configuration Application > (KAFKAMON-LOCAL/Project Name) > **Solution Package Configuration** > **Apache Kafka** > **DATA STORAGE** > **Duration** > **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.

Kafka Producers Heatmap

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](#)" display.



Filter By

Cluster Select the cluster for which you want to view data.

Fields and Data

Producers The number of producers found on the cluster and displayed in the heatmap.

Show Cluster Select this check box to display the names of the clusters at the top of each rectangle in the heatmap.

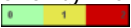




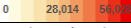
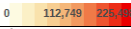
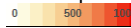
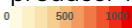
Heatmap

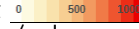
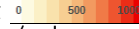
Log Scale Select this check box to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically 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.

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 Scale** is not selected.

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. 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"](#) display for a detailed view of metrics for that particular producer.

Alert Severity	<p>The current alert severity. Values range from 0 - 2, as indicated in the color gradient  bar, where 2 is the highest Alert Severity:</p> <ul style="list-style-type: none">  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	<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 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.</p>
Incoming Bytes/s	<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 KafkaProducerIncomingByteRate. The middle value in the gradient bar indicates the middle value of the range.</p> <p>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.</p>
Outgoing Bytes/s	<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 KafkaProducerOutgoingByteRate. The middle value in the gradient bar indicates the middle value of the range.</p> <p>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.</p>
Avg IO Wait Time Avg ms	<p>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.</p> <p>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.</p>
Avg Request Latency	<p>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.</p> <p>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.</p>

- 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 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.
- 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 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.

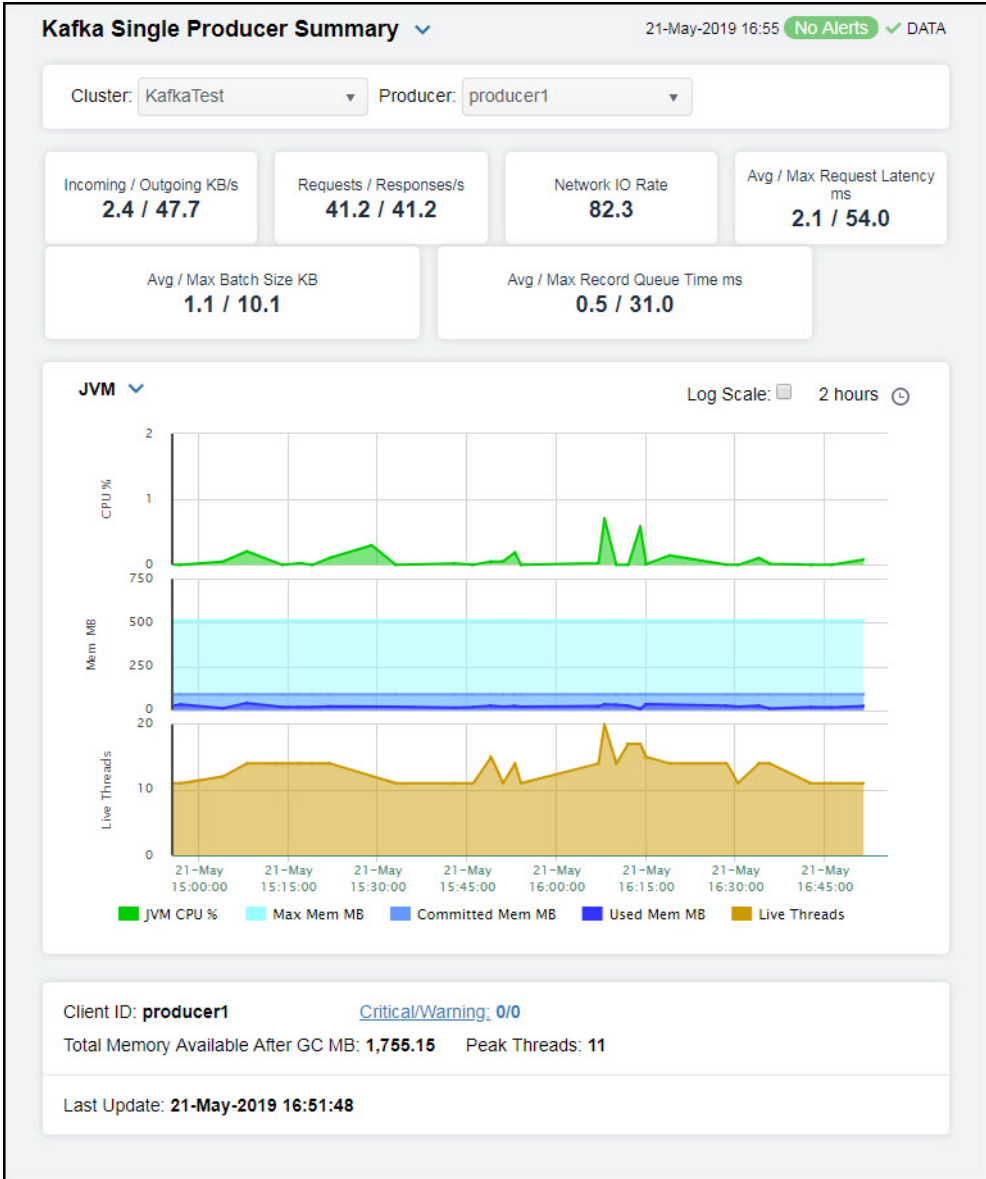
Kafka Single Producer Summary

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](#)" 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.



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.
- Producer** Select the producer for which you want to show data in the display.

Fields and Data

Incoming/ Outgoing KB/ s	The rate of incoming and outgoing bytes (kilobytes per second). *
Requests/ Responses/s	The average number of requests sent per second, and the average number of responses received per second. *
Network IO Rate	The rate of input/output network operations. *
Avg/Max Request Latency ms	The average request latency (in milliseconds), and the maximum request latency (in milliseconds). *
Avg/Max Batch Size KB	The average batch size sent by the producer, and the maximum number of messages that can be added to a batch before being sent to the event handler. *
Avg/Max Record Queue Time ms	The average time (in milliseconds) record batches spent in the record accumulator, and the maximum time (in milliseconds) record batches spent in the record accumulator. *

Trend Graphs

Performance

Requests/s -- traces the number of requests per second.

Responses/s -- traces the number of responses per second.

Max Request Latency ms -- traces the maximum request latency (in milliseconds).

Avg Request Latency ms -- traces the average request latency (in milliseconds).

Outgoing KBs/s -- traces the rate of outgoing bytes (kilobytes per second).

Avg IO Wait ms -- traces the average length of time the I/O thread spent waiting for a socket (in milliseconds).

JVM

JVM CPU % -- traces the CPU being used by the JVM.

Max Mem MB -- traces the maximum amount of available heap.

Committed Mem MB -- traces the amount of committed heap memory.

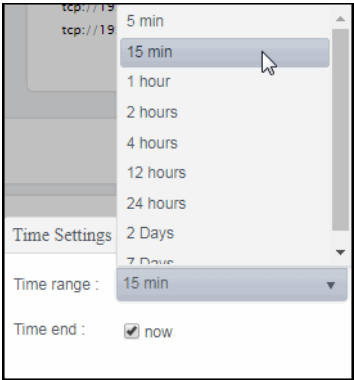
Used Mem MB -- traces the highest amount of heap used.

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.

Time Settings Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

- Client ID** The ID of the client.
- Peak Threads** The highest number of threads running at one time during the current polling period.*
- Critical/Warning** The number of critical and warning alerts.
- Total Memory Available After GC MB** The amount of memory available after garbage collection takes place.
- Last Update** The date and time of the last data update.

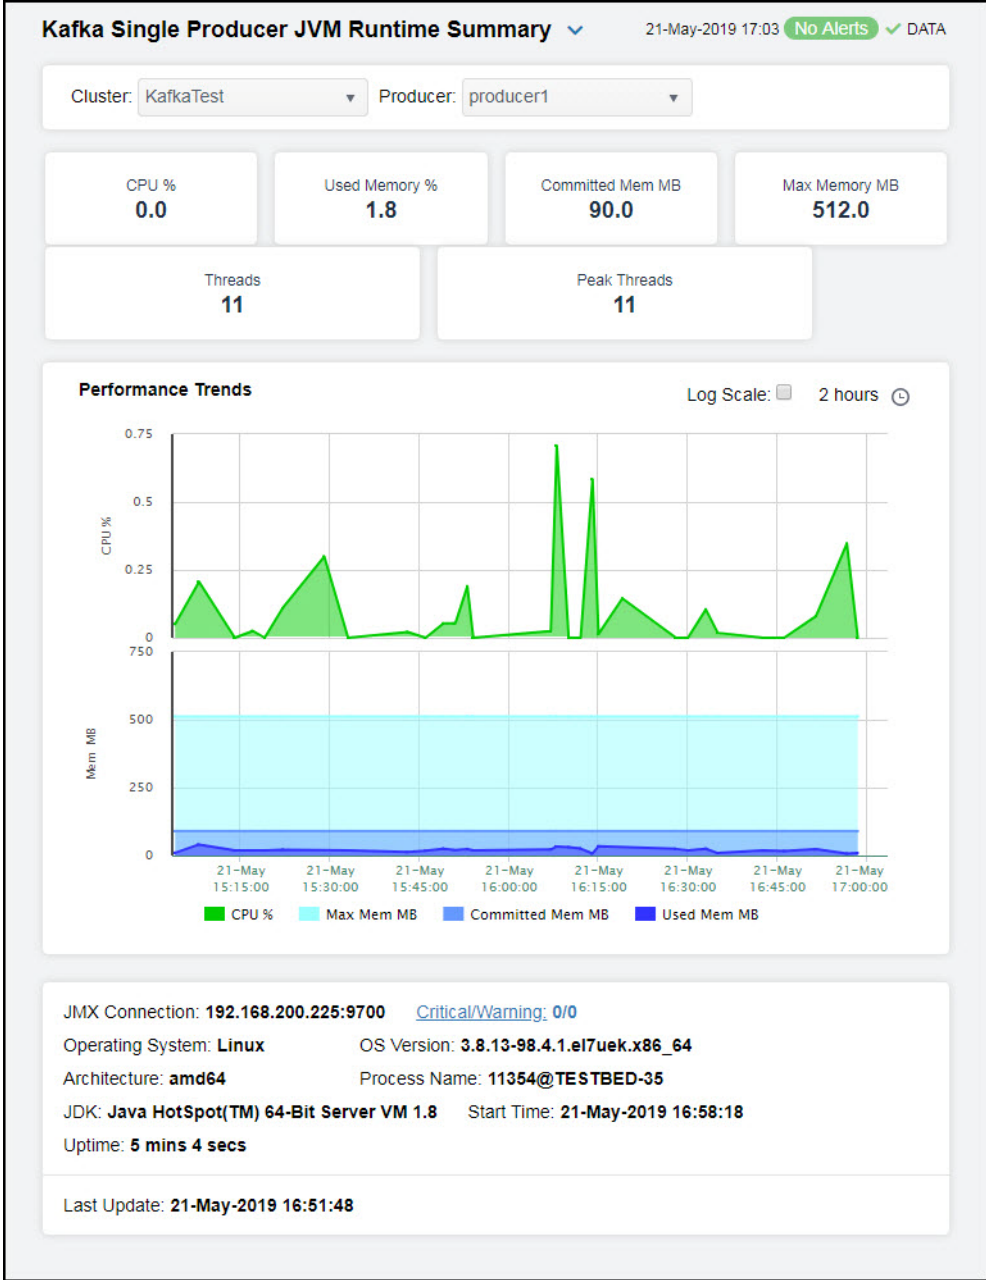
Kafka Single Producer JVM Runtime Summary

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"](#) 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.



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.

Producer	Select the producer for which you want to show data in the display.
Fields and Data	
CPU %	The percentage of CPU being used by producer.*
Used Memory %	The percentage of memory used by the producer.
Committed Mem MB	The amount of committed heap memory, in megabytes, by the producer.
Max Memory MB	The maximum amount of heap memory, in megabytes, defined for the producer.
Threads	The number of threads running on the producer.
Peak Threads	The peak number of threads running on the producer.
Performance Trends Graph	Traces the following: CPU % -- traces the percentage of CPU being used by the JVM. Max Mem MB -- traces the maximum amount of available heap, in megabytes. Committed Mem MB -- traces the amount of committed heap, in megabytes. Used Mem MB -- traces the amount of heap used, in megabytes.
JMX Connection	The name of the JMX connection.*
OS Version	The version number of the operating systems.*
Process Name	The name of the process.*
Start Time	The date and time when the producer was started.*
Critical/Warning	The number of critical and warning alerts.
Architecture	The type of processor being used.*
JDK	The JDK version number.*
Uptime	The amount of time the producer has been up and running.*
Operating System	The operating system installed on the producer.*
Last Update	The date and time of the last data update.

Kafka Consumers View

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”](#) 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”](#) 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”](#) 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”](#) 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”](#) display, which displays the lag per topic in a bar graph format and lists the lag per topic for the consumer.

Kafka 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. 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”](#) 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

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.

Consumers The number of consumers found on the selected cluster, which are listed in the **All Consumers Table**.

Table

Cluster The name of the cluster.

Consumer 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.

Alert Count The total number of alerts for the host.

Client ID The ID of the consumer.*

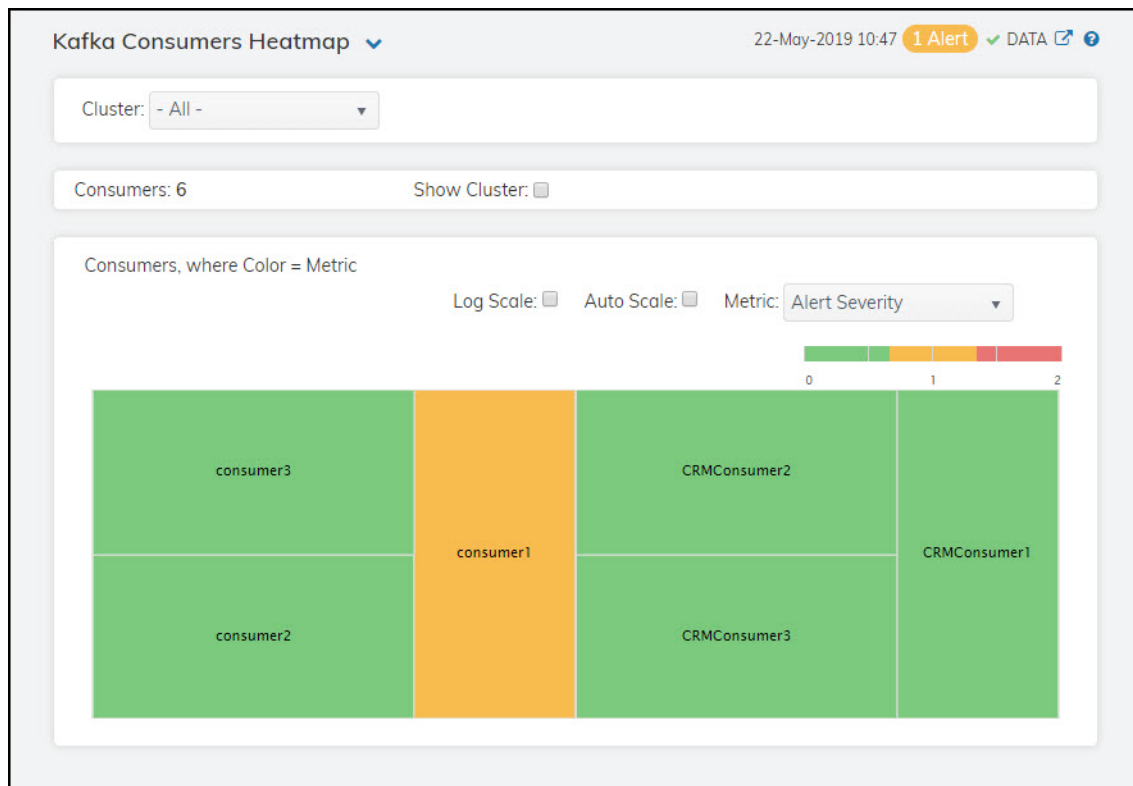
Consumed Bytes/s	The average number of bytes consumed per second.*
Avg Fetch Latency	The average time taken for a fetch request.*
Max Fetch Latency	The maximum time taken for a fetch request.*
Fetch Records/s	The number of fetch requests per second.*
Fetch Size Avg bytes	The average number of bytes fetched per request.*
Max Fetch Size bytes	The maximum number of bytes fetched per request.*
Fetch Throttle Avg Time ms	The average throttle time in milliseconds.*
Fetch Throttle Max Time ms	The maximum throttle time in milliseconds.*
Consumed Records/s	The average number of records consumed per second.*
Lag	The difference between the current consumer position in the partition and the end of the log.*
Lag Rate	The rate of consumer lag.
Records per Request Avg	The average number of records in each request.*
JMX Connection String	The JMX connection string.*
Kafka Version	The current version of Apache Kafka installed.*
Connected	Denotes whether or not the consumer is connected.*
Expired	<p>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.</p> <p>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.</p>
Timestamp	The date and time the row data was last updated.

Kafka Consumers Heatmap

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"](#) display.



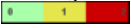






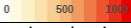
Filter By:

Cluster Select the cluster for which you want to view data.


Consumers The number of consumers found on the cluster and that are displayed in the heatmap.

Show Cluster Select this check box to display the names of the clusters at the top of each rectangle in the heatmap.

Heatmap

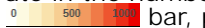
Log Scale	Select this check box to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not 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.
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 Scale is not selected.
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. 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 " 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: <ul style="list-style-type: none">  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 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.
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 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.
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 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.

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 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.

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 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.

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 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.

Kafka Single Consumer Summary

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"](#) 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.



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.

Fields and Data

- Consumed Bytes/s** The average number of bytes consumed per second. *

Consumed Records/s	The average number of records consumed per second.*
Avg Records/Request	The average number of records in each request.*
Avg Fetch Records/s	The average number of fetch requests per second.*
Avg/Max Fetch Latency	The average time taken for a fetch request, and the maximum amount of time taken for a fetch request.*
Consumer Lag	The maximum lag in number of records for any partition.*

Trend Graphs**Performance**

JVM CPU % -- traces the CPU percentage being used by the JVM.

Max Record Lag -- traces the maximum record lag.

Consumer Lag -- traces the lag in number of records for any partition.

Fetch Rate -- traces the fetch rate.

Consumption Rate

Consumed KB/s -- traces the rate of kilobytes being consumed, per second.

Consumed Records/s -- traces the rate of records being consumed, per second.

JVM Resource Utilization

JVM CPU % -- traces the CPU percentage being used by the JVM.

Max Mem MB -- traces the maximum amount of available heap, in megabytes.


Committed Mem MB -- traces the amount of committed heap, in megabytes.

Used Mem MB -- 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.

Time Settings Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

Records Lag Max	The maximum lag in number of records for any partition. *
Fetch Throttle Max Time ms	The maximum throttle time, in milliseconds. *
Critical/Warning	The number of critical and warning alerts.
Fetch Throttle Avg Time ms	The average throttle time, in milliseconds. *
Fetch Records/s	The average number of fetch requests per second. *
Fetch Size Avg bytes	The average number of bytes fetched per request. *
Max Fetch Size bytes	The highest number of bytes fetched per request. *
Last Update	The date and time of the last data update.

Kafka Single Consumer JVM Runtime Summary

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"](#) 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.



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.

Fields and Data

- CPU %** The percentage of CPU used by the consumer.*

Used Memory %	The percentage of memory used of this consumer as JVM.
Committed Mem MB	The committed heap memory, in megabytes, of this consumer as JVM.
Max Memory MB	The maximum heap memory, in megabytes, defined for this consumer.
Threads	The number of threads running on this consumer.
Peak Threads	The peak number of threads running on this consumer.
Performance Trends Graph	<p>Traces the following:</p> <ul style="list-style-type: none"> CPU % -- traces the CPU percentage. Max Mem MB -- traces the maximum amount of available heap, in megabytes. Committed Mem MB -- traces the amount of committed heap, in megabytes. Used Mem MB -- traces the highest amount of heap used, in megabytes. <p>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.</p>

Time Settings Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



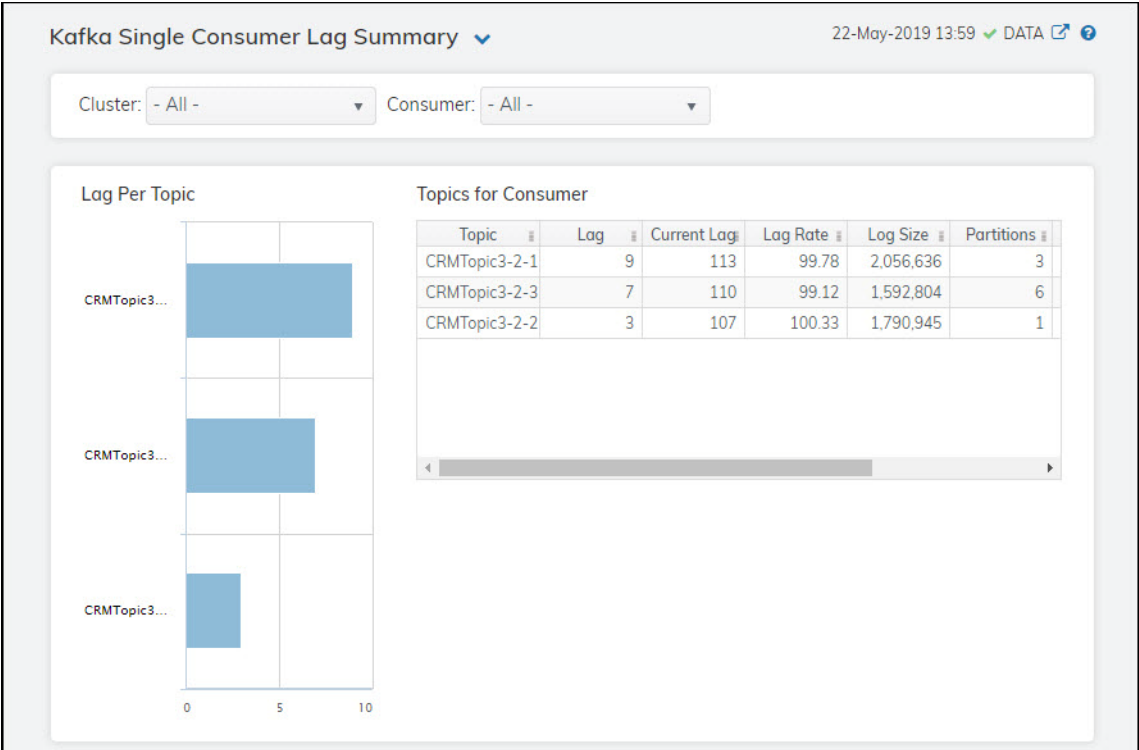
To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

JMX Connection	The name of the JMX connection.*
OS Version	The version number of the operating systems.*
Process Name	The name of the process.*
Start Time	The date and time when the producer was started.*
Critical/Warning	The number of critical and warning alerts.
Architecture	The type of processor being used.*
JDK	The JDK version number.*
Uptime	The amount of time the producer has been up and running.*
Operating System	The operating system installed on the producer.*
Last Update	The date and time of the last data update.

Kafka Single Consumer Lag Summary

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"](#) 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.



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.

Lag Per Topic Bar Graph Displays the lag per topic in a bar graph format.

Topics for Consumer Table

Topic The name of the topic.

Lag The difference between the current consumer position in the partition and the end of the log.*

Current Lag The difference in the amount of lag from the previous polling period to the current polling period.*

Lag Rate The rate of change in the amount of lag.*

Log Size	The current number of messages in the log.*
Partitions	The number of partitions containing the topic.
Time Stamp	The date and time the row data was last updated.

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.

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 Database"](#)
- ["Connector for Oracle Enterprise Manager"](#)
- ["Oracle WebLogic"](#)

Note: This document assumes familiarity with the products monitored. For additional details, refer to vendor documentation.

Oracle Coherence

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"](#) (pictured below), and several Views which can be found under **Components** tab > **Databases** > **Oracle Database**:

- ["Oracle Database Overview Display"](#): This display presents a health snap-shot of the Oracle Database system.
- ["Oracle Databases View"](#): These displays present metrics about Oracle databases.
- ["Oracle Instances View"](#): These displays present metrics about instances on Oracle databases.

Oracle Database Overview Display

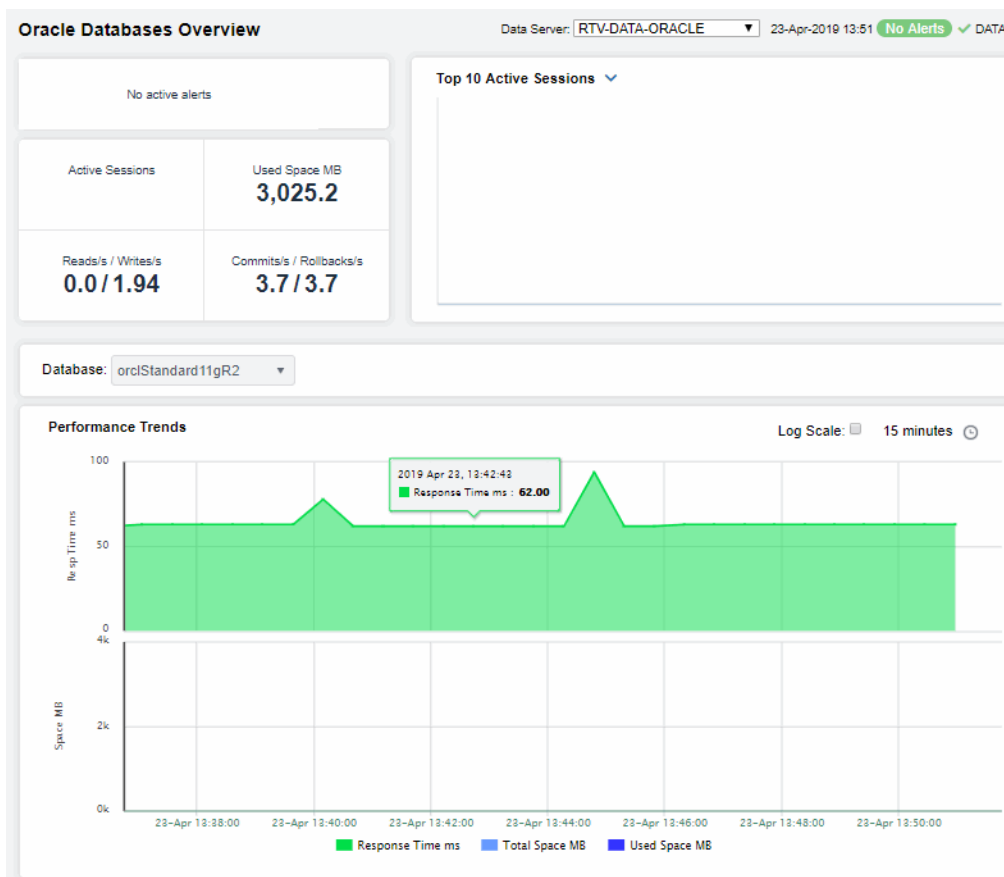
The **Oracle Database Overview** is the top-level display for the Oracle Database Solution Package, which provides a good starting point for immediately getting the status of all your Oracle Databases on your Data Server.

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"](#) 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

These displays present performance metrics for Oracle databases. Displays are:

- [“Oracle Databases Table”](#): A list of all Oracle databases with detailed utilization metrics.
- [“Databases Heatmap”](#): A heatmap shows alert status of all Oracle databases.
- [“Databases Summary”](#): Detailed performance metrics for one Oracle database.




Oracle Databases Table

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”](#) 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

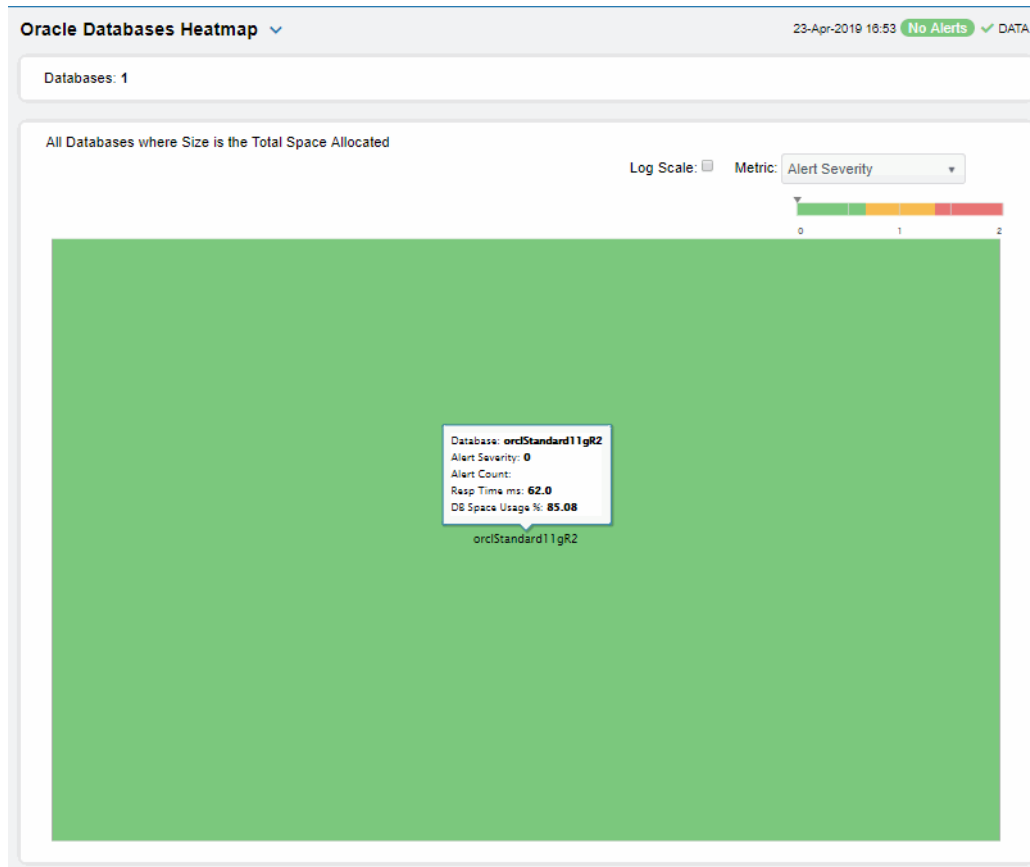
View status and alerts for all Oracle databases. Use this display to quickly identify a database with performance or utilization issues.

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"](#) 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

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](#)” 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

These displays present performance metrics for Oracle Database instances. Displays are:

- [“Oracle Instances Table”](#): A list of all Oracle Database instances with detailed utilization metrics.
- [“Instances Heatmap”](#): A heatmap shows alert status of Oracle Database instances.
- [“Instance Summary”](#): Detailed performance metrics for one Oracle Database instance.

Oracle Instances Table

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”](#) 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.

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

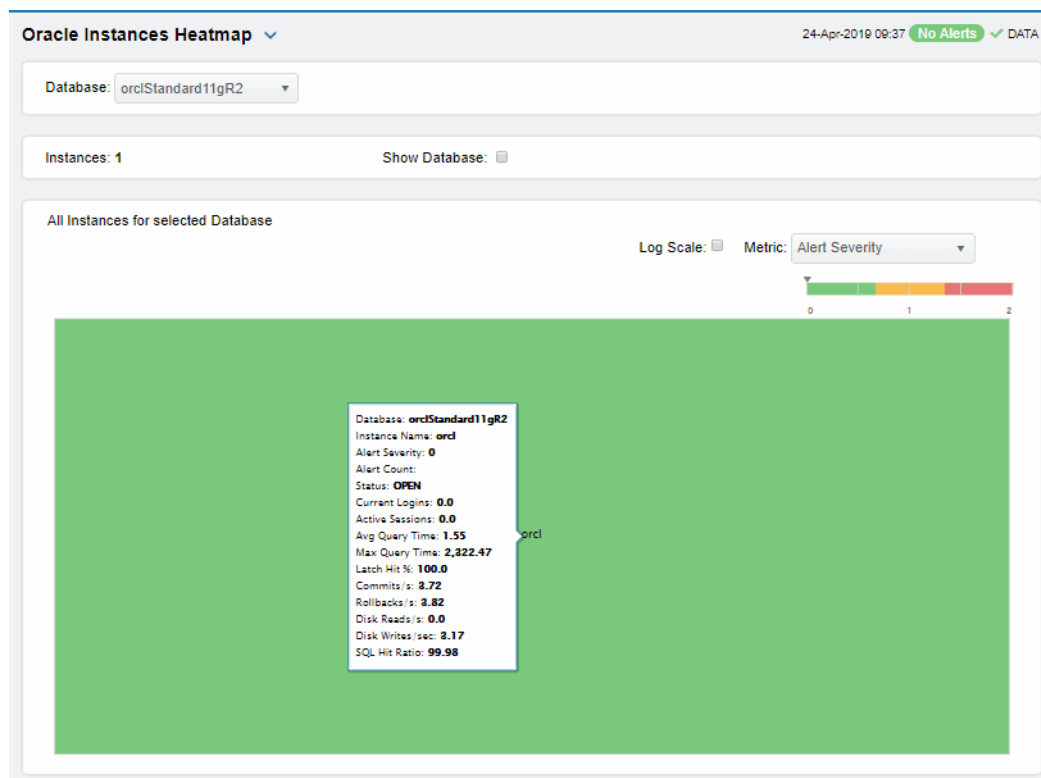
View status and alerts for all instances on one or all Oracle databases. Use this display to quickly identify an instance with performance or utilization issues.

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”](#) 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:

Alert Severity	<p>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.</p> <ul style="list-style-type: none">  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.
Current Logins	<p>The number of users logged on. The color gradient  bar values range from 0 to the maximum count in the heatmap. The middle value in the gradient bar indicates the average count.</p>
Active Sessions	<p>The number of active sessions. The color gradient  bar values range from 0 to the maximum count in the heatmap. The middle value in the gradient bar indicates the average count.</p>
Avg Query Time	<p>The average amount of time, in seconds, to perform a query. The color gradient  bar values range from 0 to the maximum average in the heatmap. The middle value in the gradient bar indicates the average amount.</p>
Max Query Time	<p>The maximum amount of time, in seconds, to perform a query. The color gradient  bar values range from 0 to the maximum amount in the heatmap. The middle value in the gradient bar indicates the average maximum amount.</p>
Latch Hit Ratio	<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>
Data Dict Hit Ratio	<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>
SQL Hit Ratio	<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>
Commits/sec	<p>The number of commits per second. The color gradient  bar values range from 0 to the maximum count in the heatmap. The middle value in the gradient bar indicates the average.</p>
Rollbacks/sec	<p>The number of rollbacks per second. The color gradient  bar values range from 0 to the maximum count in the heatmap. The middle value in the gradient bar indicates the average.</p>
Disk Reads/sec	<p>The number of disk reads per second. The color gradient  bar values range from 0 to the maximum count in the heatmap. The middle value in the gradient bar indicates the average.</p>
Disk Writes/sec	<p>The number of disk writes per second. The color gradient  bar values range from 0 to the maximum count in the heatmap. The middle value in the gradient bar indicates the average.</p>


Instance Summary

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](#)” 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

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”](#) (pictured below), and several Views which can be found under **Components** tab > **Databases** > **Oracle Database**:

- [“Oracle Database Overview Display”](#): This display presents a health snap-shot of the Oracle Database system.
- [“Oracle Databases View”](#): These displays present metrics about Oracle databases.
- [“Oracle Instances View”](#): These displays present metrics about instances on Oracle databases.Si

Oracle Database Overview Display

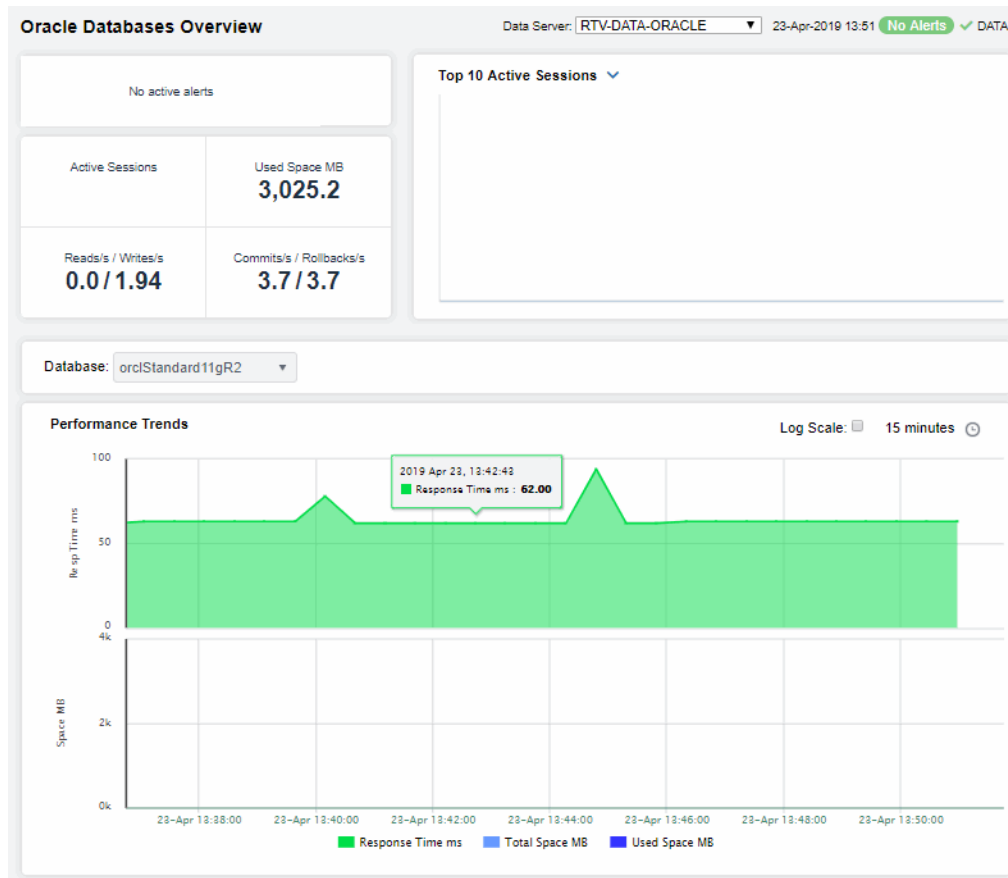
The **Oracle Database Overview** is the top-level display for the Oracle Database Solution Package, which provides a good starting point for immediately getting the status of all your Oracle Databases on your Data Server.

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”](#) 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

These displays present performance metrics for Oracle databases. Displays are:

- **"Oracle Databases Table"**: A list of all Oracle databases with detailed utilization metrics.
- **"Databases Heatmap"**: A heatmap shows alert status of all Oracle databases.
- **"Databases Summary"**: Detailed performance metrics for one Oracle database.

Oracle Databases Table

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](#)” 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
orclStandard11gR2	✔		✔		1	63.0	85.08	3,021.25	542.8

Databases Heatmap

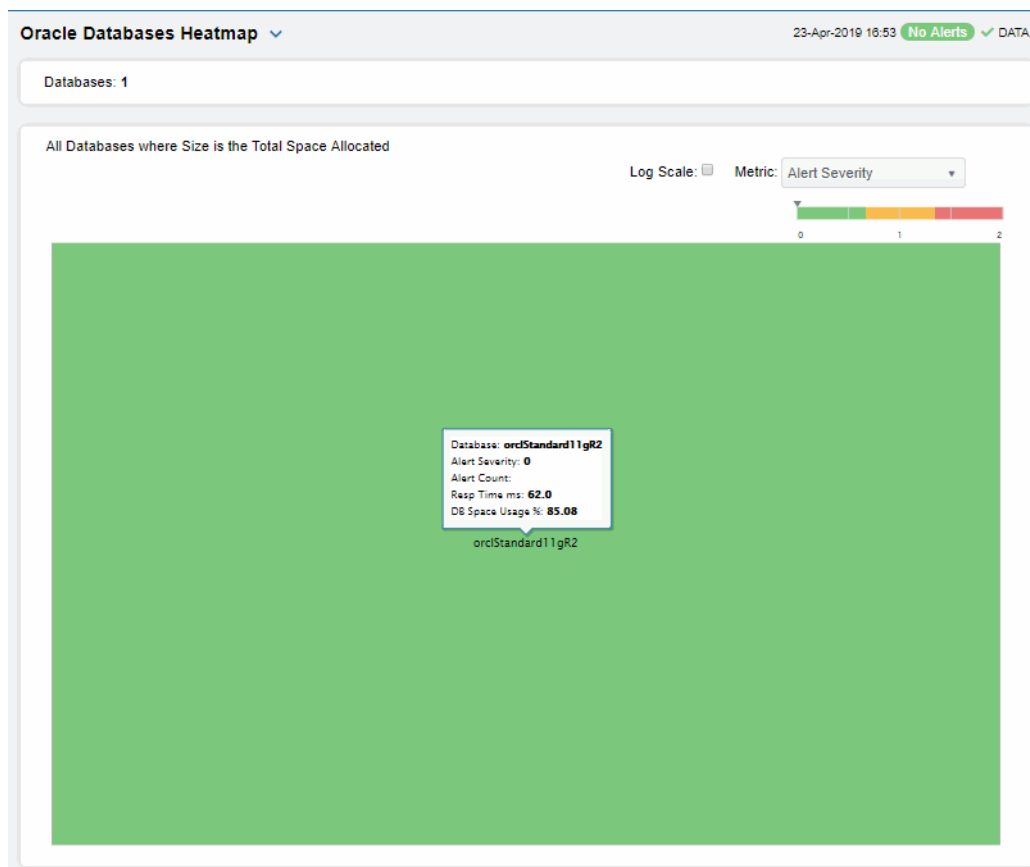
View status and alerts for all Oracle databases. Use this display to quickly identify a database with performance or utilization issues.

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”](#) 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:

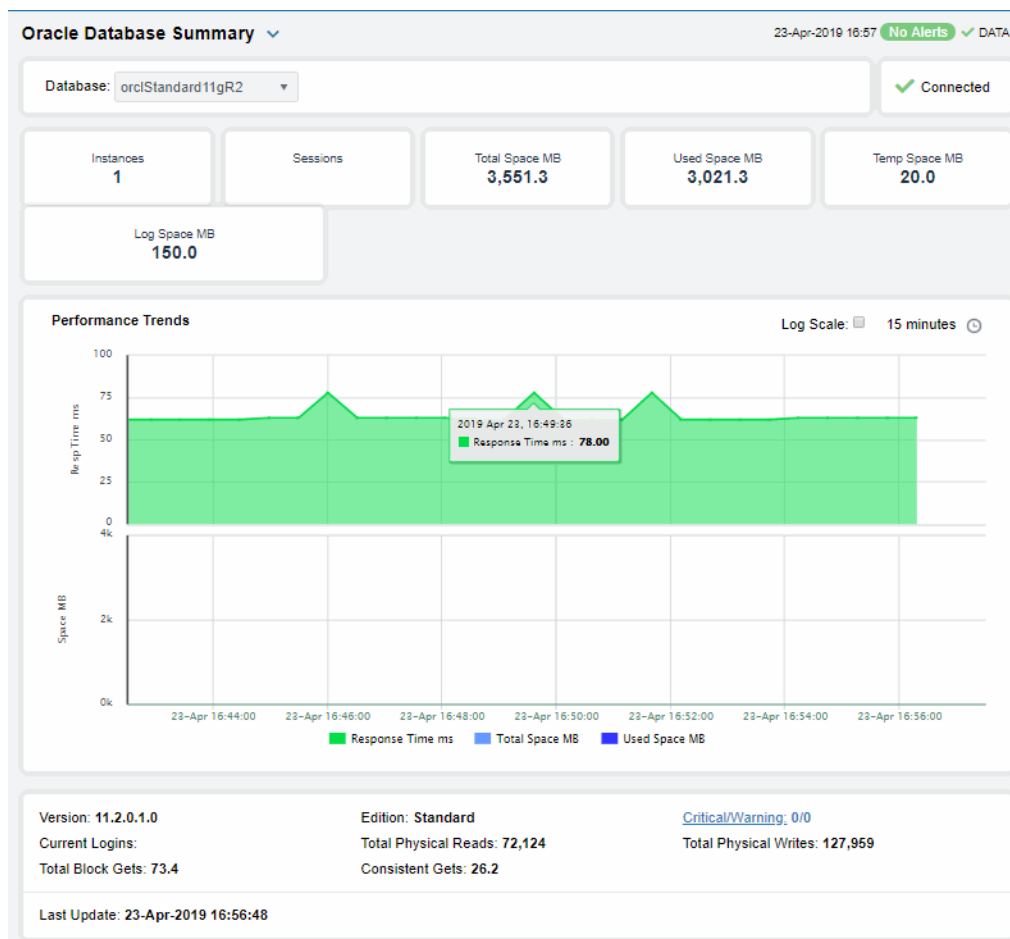
Alert Severity	<p>The maximum level of alerts in the heatmap rectangle. Values range from 0 - 2, as indicated in the color gradient  bar, where 2 is the highest Alert Severity.</p> <ul style="list-style-type: none">  Red indicates that one or more metrics have reached their alarm threshold. Metrics that 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	<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 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.</p>
Response Time	<p>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.</p>
Database Space Usage	<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 0 to the maximum amount of space used in the heatmap. The middle value in the gradient bar indicates the average amount.</p>

Databases Summary

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](#)” 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

These displays present performance metrics for Oracle Database instances. Displays are:

- [“Oracle Instances Table”](#): A list of all Oracle Database instances with detailed utilization metrics.
- [“Instances Heatmap”](#): A heatmap shows alert status of Oracle Database instances.
- [“Instance Summary”](#): Detailed performance metrics for one Oracle Database instance.

Oracle Instances Table

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”](#) 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.

Database	Instance Id	Instance Name	Expired	Alert Level	Alert Count	Instance Role	DB Status	Version	Active Sessions	Ina
ordStandard11gR2	0	orcl		✓		PRIMARY_INSTAN	ACTIVE	11.2.0.1.0	0.0	

Instances Heatmap


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](#)” 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:

Alert Severity	<p>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.</p> <ul style="list-style-type: none">  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.
Current Logins	<p>The number of users logged on. The color gradient  bar values range from 0 to the maximum count in the heatmap. The middle value in the gradient bar indicates the average count.</p>
Active Sessions	<p>The number of active sessions. The color gradient  bar values range from 0 to the maximum count in the heatmap. The middle value in the gradient bar indicates the average count.</p>
Avg Query Time	<p>The average amount of time, in seconds, to perform a query. The color gradient  bar values range from 0 to the maximum average in the heatmap. The middle value in the gradient bar indicates the average amount.</p>
Max Query Time	<p>The maximum amount of time, in seconds, to perform a query. The color gradient  bar values range from 0 to the maximum amount in the heatmap. The middle value in the gradient bar indicates the average maximum amount.</p>
Latch Hit Ratio	<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>
Data Dict Hit Ratio	<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>
SQL Hit Ratio	<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>
Commits/sec	<p>The number of commits per second. The color gradient  bar values range from 0 to the maximum count in the heatmap. The middle value in the gradient bar indicates the average.</p>
Rollbacks/sec	<p>The number of rollbacks per second. The color gradient  bar values range from 0 to the maximum count in the heatmap. The middle value in the gradient bar indicates the average.</p>
Disk Reads/sec	<p>The number of disk reads per second. The color gradient  bar values range from 0 to the maximum count in the heatmap. The middle value in the gradient bar indicates the average.</p>
Disk Writes/sec	<p>The number of disk writes per second. The color gradient  bar values range from 0 to the maximum count in the heatmap. The middle value in the gradient bar indicates the average.</p>


Instance Summary

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](#)” 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

The HTML version features an overview display, [“WebLogic Overview Display”](#) (pictured below), and the following Views which can be found under **Components** tab > **Application/ Web Servers** > **Oracle WebLogic**:

- [“WebLogic Servers View”](#)
- [“Single WebLogic Server View”](#)
- [“WebLogic Applications View”](#)

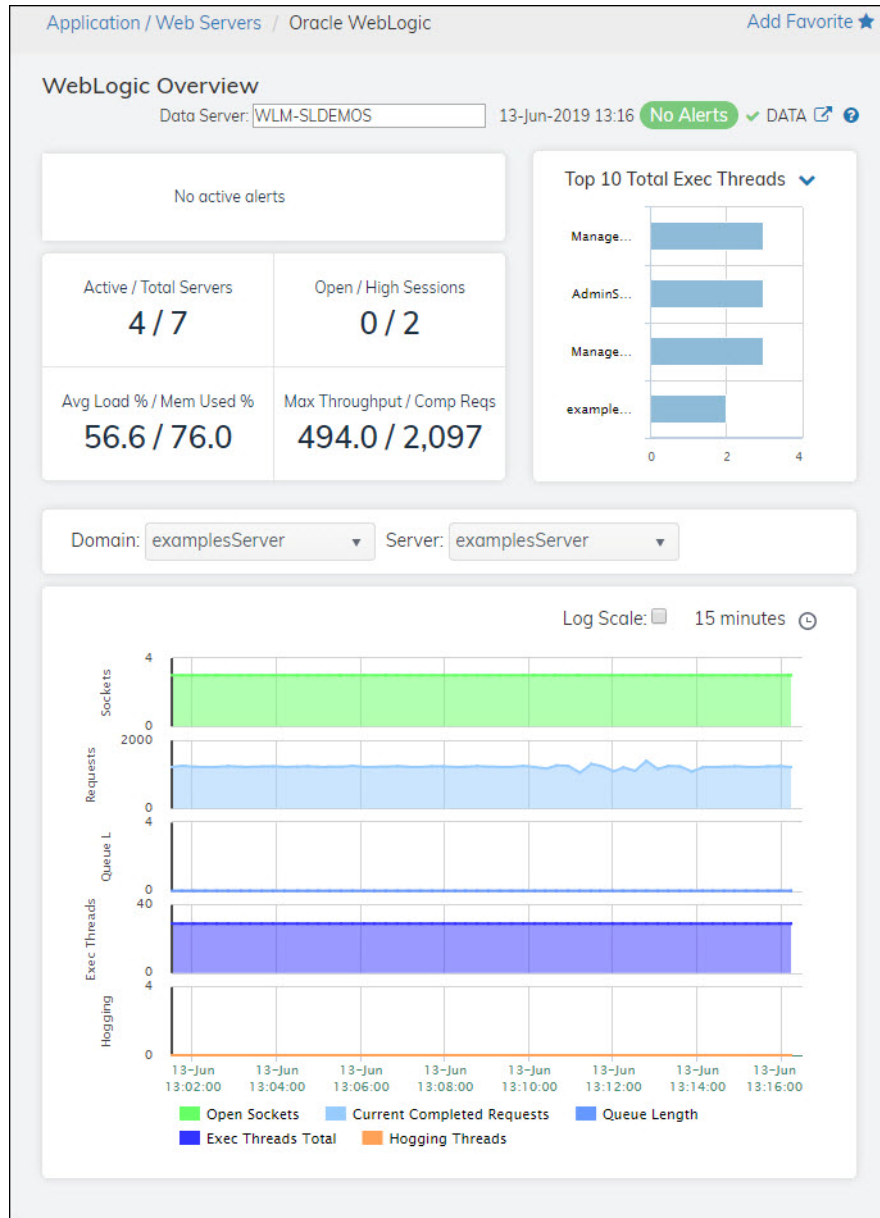
WebLogic Overview Display

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

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](#)” 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](#)” display, which shows status and alerts for all Oracle® WebLogic servers in a heatmap.
- **Clusters Table:** Opens the “[WebLogic Clusters Table](#)” display, which enables you to track utilization and performance metrics for all clusters on a particular domain, or on all domains.

WebLogic 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. Double-click on a table row to drill-down to the “[WebLogic Server Summary](#)” 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

Filter By

Domain Select the domain for which you want to view data, or select **All Domains** to view data for all domains.

Fields and Data

Servers The total number of servers on the cluster.




Active The total number of active servers on the cluster.

Show Active Servers Only

Select this check box to display only the active servers in the table.

Table

This table shows information for the selected domain/cluster(s) combination. Click on a table row to drill-down to the ["WebLogic Server Summary"](#) display and view metrics for that particular server.

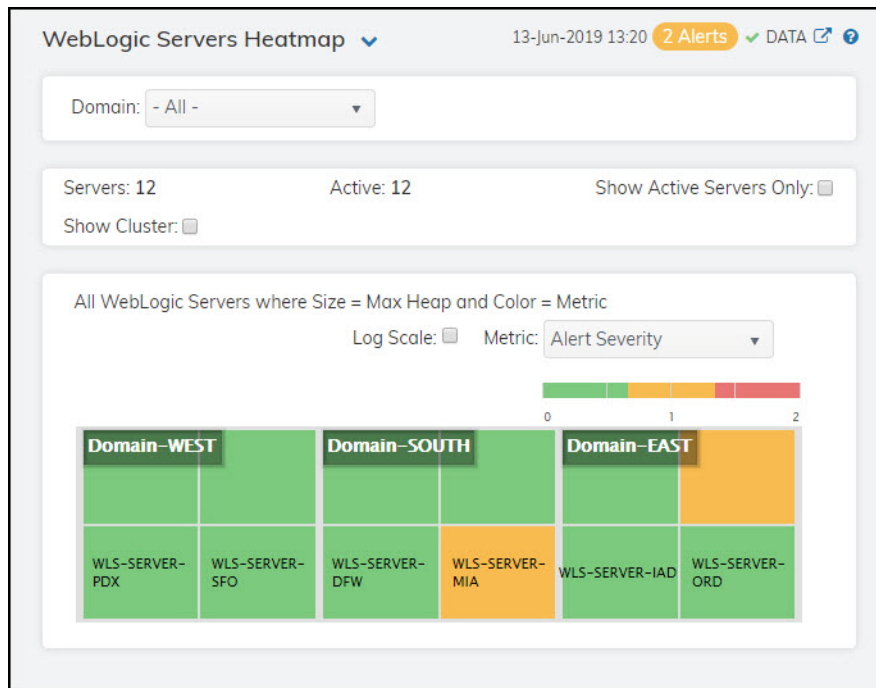
Domain	The name of the domain.
Cluster	The name of the cluster.
Server	The name of the server.
State	The current state of the server.
Alert Level	The current alert level.  -- One or more alerts have exceeded their specified ALARMLEVEL threshold.  -- One or more alerts have exceeded their specified WARNINGLEVEL threshold.  -- No alerts have exceeded an alert threshold.
Port	The port on which this server is listening for SSL connections.*
JVM Proc Load %	A snapshot of the load that the virtual machine is placing on all processors on the host computer.*
Heap Free %	The percentage of free heap memory on the server.*
HeapSizeMax	The maximum amount of heap, in bytes, available for use.*
Used Heap	The total amount of heap used, in bytes.*
Current Heap	The current size of the JVM heap, in bytes, being used.*
Open Sockets	The current number of sockets registered for socket muxing on this server.*
Hogging Threads	The number of hogging threads on the server.*
Exec Threads Total	The current number of execute threads.*
Exec Idle Threads	The current number of idle threads.
Restarts	The total number of restarts for this server since the cluster was last started.*
All Proc Avg Load	The average load percentage for all processors on the host computer.*
Shutting Down	When checked, denotes that the server is currently shutting down.*
Restart Required	When checked, denotes that the server needs to be restarted in order to activate configuration changes.*
Uptime	The length of time (in milliseconds) that the server has been up and running.*
Startup Time	The length of time (in milliseconds) that it took for the server to start up.*

WebLogic Version	The current version of WebLogic running on the server.*
JVM Type	The type of JVM currently being used on the server.*
Java Version	The current version of Java running on the server.*
JavaVendor	The name of the vendor of the Java version running on the server.*
OS	The name of the operating system running on the server.*
Expired	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.
Time Stamp	The date and time this row of data was last updated.

WebLogic Servers Heatmap

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](#)" 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.

**Filter By:****Domain**

Select the domain (or **All Domains**) from the drop down list for which you want to view data.

Fields and Data**Servers**

The total number of active, inactive, and standby servers.

Active

The number of active servers listed in the display.

Show Inactive Servers Only

Select this check box to display only the inactive servers in the heatmap.

Show Cluster

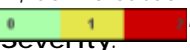





Select this check box to display the names of the clusters in the heatmap

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

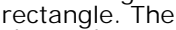
Select the metric driving the heatmap display. The default is Alert Severity. Each **Metric** has a color gradient bar that maps values to colors. The heatmap organizes the instances by host, where each rectangle represents an instance. Mouse-over any rectangle to display the current values of the metrics for the instance. Click on a rectangle to drill-down to the associated ["WebLogic Server Applications Summary"](#) display for a detailed view of metrics for that particular domain.

Alert Severity	<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 2 is the greatest Alert Severity.</p> <p>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.</p> <p>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.</p> <p>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.</p>
Alert Count	<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 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>
Jvm CPU %	<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 0 to the alert threshold of WisServerCpuHigh. The middle value in the gradient bar indicates the middle value of the range.</p>
Host CPU %	<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 0 to the alert threshold of WisServerHostCpuHigh. The middle value in the gradient bar indicates the middle value of the range.</p>
JVM Memory %	<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 0 to the alert threshold of WisMemoryUsageHigh. The middle value in the gradient bar indicates the middle value of the range.</p>
Open Sockets	<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 0 to the alert threshold of WisOpenSocketsHigh. The middle value in the gradient bar indicates the middle value of the range.</p>

Thread Total Count

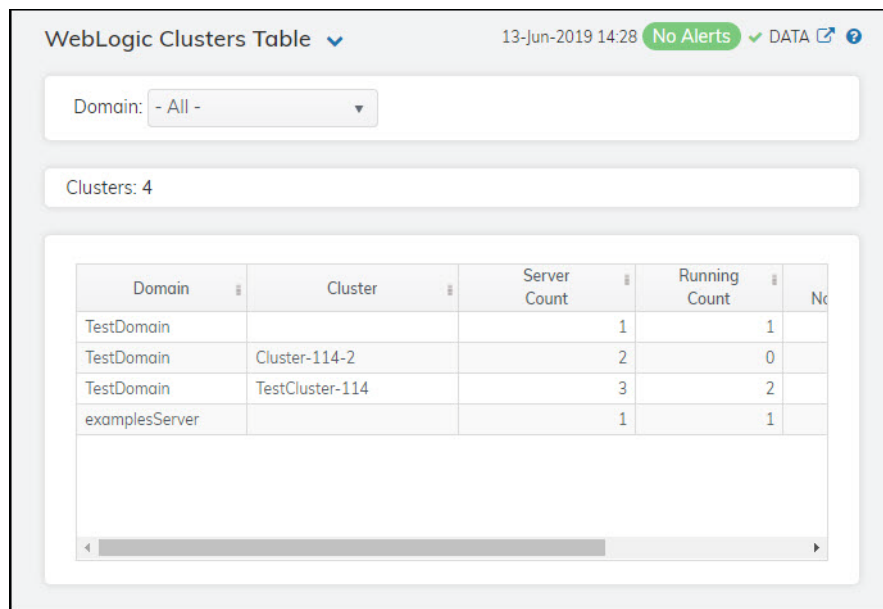
The total number of threads in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the alert threshold of **WisThreadsTotalHigh**. The middle value in the gradient bar indicates the middle value of the range.

Hogging Threads

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 0 to the alert threshold of **WisHoggingThreadsHigh**. The middle value in the gradient bar indicates the middle value of the range.

WebLogic Clusters Table

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**” 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	No
TestDomain		1	1	
TestDomain	Cluster-114-2	2	0	
TestDomain	TestCluster-114	3	2	
examplesServer		1	1	

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.

Filter By

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.

Clusters The current number of clusters listed in the table.

Table

Lists the clusters in the currently selected domain, or lists all clusters in all domains.

Domain	The name of the domain
Cluster	The name of the cluster.*
Server Count	The total number of servers on the cluster.*
Running Count	The total number of servers running on the cluster.*
Percent Not Running	The percentage of servers not running on the cluster.*
Time Stamp	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. Clicking **Single WebLogic Server** from the left menu opens the [“WebLogic Server Summary”](#) 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”](#) display, which displays the JVM details for a particular WebLogic server on a specific domain.
- **Server Detail:** Opens the [“WebLogic Server Detail Tables”](#) 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”](#) display, which displays JDBC module utilization, performance, and trend data for a specific WebLogic server.
- **Work Manager:** Opens the [“WebLogic Server Work Manager Table”](#) 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”](#) display, which displays available utilization and performance data for all configurations on a specific domain.

WebLogic Server Summary

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](#)” 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.



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

- Connection** Select the connection for which you want to view data.
- Location** Select the name of the WLS Server located on the specified connection for which you want to view data.

Fields and Data

- Open Sockets** The number of current open sockets for the server.
- Current Completed Requests** The number of current completed requests on the server.
- Hogging Threads** The current number of hogging threads. *
- Thread Queue Length** The current thread queue length. *
- Execute Threads** The current number of execute threads. *
- Total Restarts** The total number of times the server has restarted since the last update time. *

Trend Graph

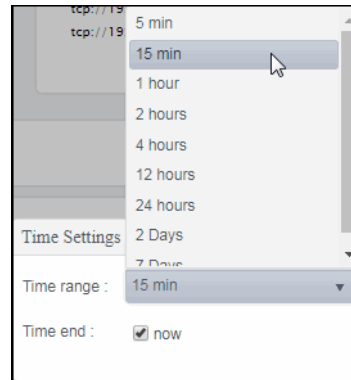
Displays data and a trend graph for the following:

- Open Sockets**-- Traces the number of open sockets of the server.
- Completed Requests**-- Traces the number of completed requests on the server.
- Queue Length**-- Traces the queue length on the server.
- Exec 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.

Time Settings

Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

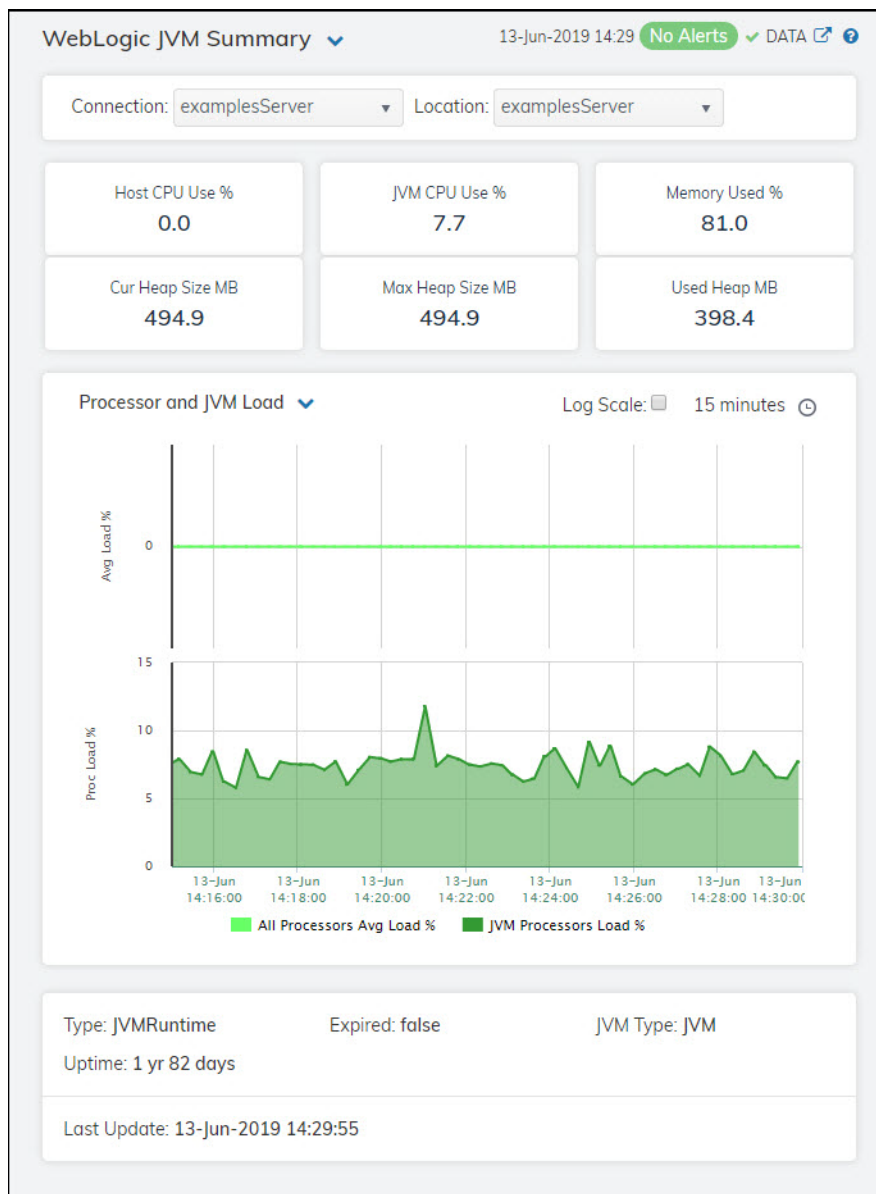
Cluster	The name of the cluster.
Port	The name of the port.*
Activation Time	The date and time in which the server was started.*
Admin Host	The IP address of the administration server's host.*
Expired	When true , the data has exceeded the specified cache expiration time (set by default at 45 seconds) and is no longer current. Once the cache has been refreshed and is displaying current data, this field will be set to false . This field will be set to false as long as the cache has been refreshed within the specified cache expiration time and the data is current.
Listen Port Enabled	When true , denotes that the listen port is enabled.*
Cluster Master	When true , denotes that the server is a cluster master.*
Admin Port	The name of the administration server's port.*
Host	The name of the host.*
State	The current state of the WebLogic server.*

Admin Server	Indicates whether the server is an Administration Server.*
Admin Port Enabled	When true, denotes that the admin port is enabled.*
Last Update	The date and time of the last data update.

WebLogic JVM Summary

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](#)” 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.



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.

Filter By:

Connection	Select the connection for which you want to view data.
Location	Select the name of the WLS Server located on the specified connection for which you want to view data.

Fields and Data

Host CPU Use %	The current CPU usage percentage on the host.*
JVM CPU Use %	The current JVM CPU usage percentage.*
Memory Used %	The percentage of memory used for the server.
Curr Heap Size MB	The current size of the JVM heap, in megabytes.*
Max Heap Size MB	The maximum amount of available heap memory, in megabytes.*
Used Heap MB	The amount of heap memory used, in megabytes.*

Trend Graph**Processor and JVM Load**

All Processors Avg Load % -- Traces the percentage of the host CPU being used.

JVM Processors Load % -- Traces the percentage of the JVM CPU being used.

Heap Memory

Max Heap MB -- Traces the maximum amount of heap memory available (in megabytes).

Cur Heap MB -- Traces the current size of the JVM heap, in megabytes.


Used Heap MB -- Traces the total amount of heap used, 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.

Time Settings

Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

Type	The type of JVM currently being used on the server.*
Uptime	The amount of time since the server was last started, shown in days, hours, and minutes (for example, 1d 23:43).*
Expired	When true , 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, this field will be set to false . This field will be set to false as long as the cache has been refreshed within the specified cache expiration time and the data is current.
JVM Type	The type of JVM currently being used on the server.*
Last Update	The date and time of the last data update.

WebLogic Server Detail Tables

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](#)” 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 Tota
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

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.

Filter By:

- Connection** Select the connection for which you want to view data.
- Location** Select the name of the WLS Server located on the specified connection for which you want to see data.

Server Runtime Table

Domain	The name of the connection.
Server	The name of the WLS Server located on the specified connection.
Cluster	The name of the cluster.
Activation Time	The time when the server was started.*
Admin Server	Indicates whether the server is an Administration Server.*
Admin Host	The address on which the Administration Server is listening for connections.*
AdminServerListenPort	The port on which the Administration Server is listening for connections.*
AdminServerListenPort Secure	Indicates whether the port that the server uses for administrative traffic is configured to use a secure protocol.*
Admin Port	The port on which this server is listening for administrative requests.*
Admin Port Enabled	Indicates whether the administration port is enabled on the server.*
Cluster Master	When checked, denotes that the cluster is a cluster master.*
Current Directory	The absolute path of the directory from which the server was started.*
Host	The machine on which the server is running.*
Default URL	The URL that clients use to connect to this server's default network channel.*
Health State	The health state of the server as reported by the server's self-health monitoring.*
Listen Address	The address on which this server is listening for connections through the default network channel.*
Port	The port on which this server is listening for connections.*
Listen Port Enabled	Indicates whether the default listen port is enabled on the server.*
Name	The name of the Java Virtual Machine.*

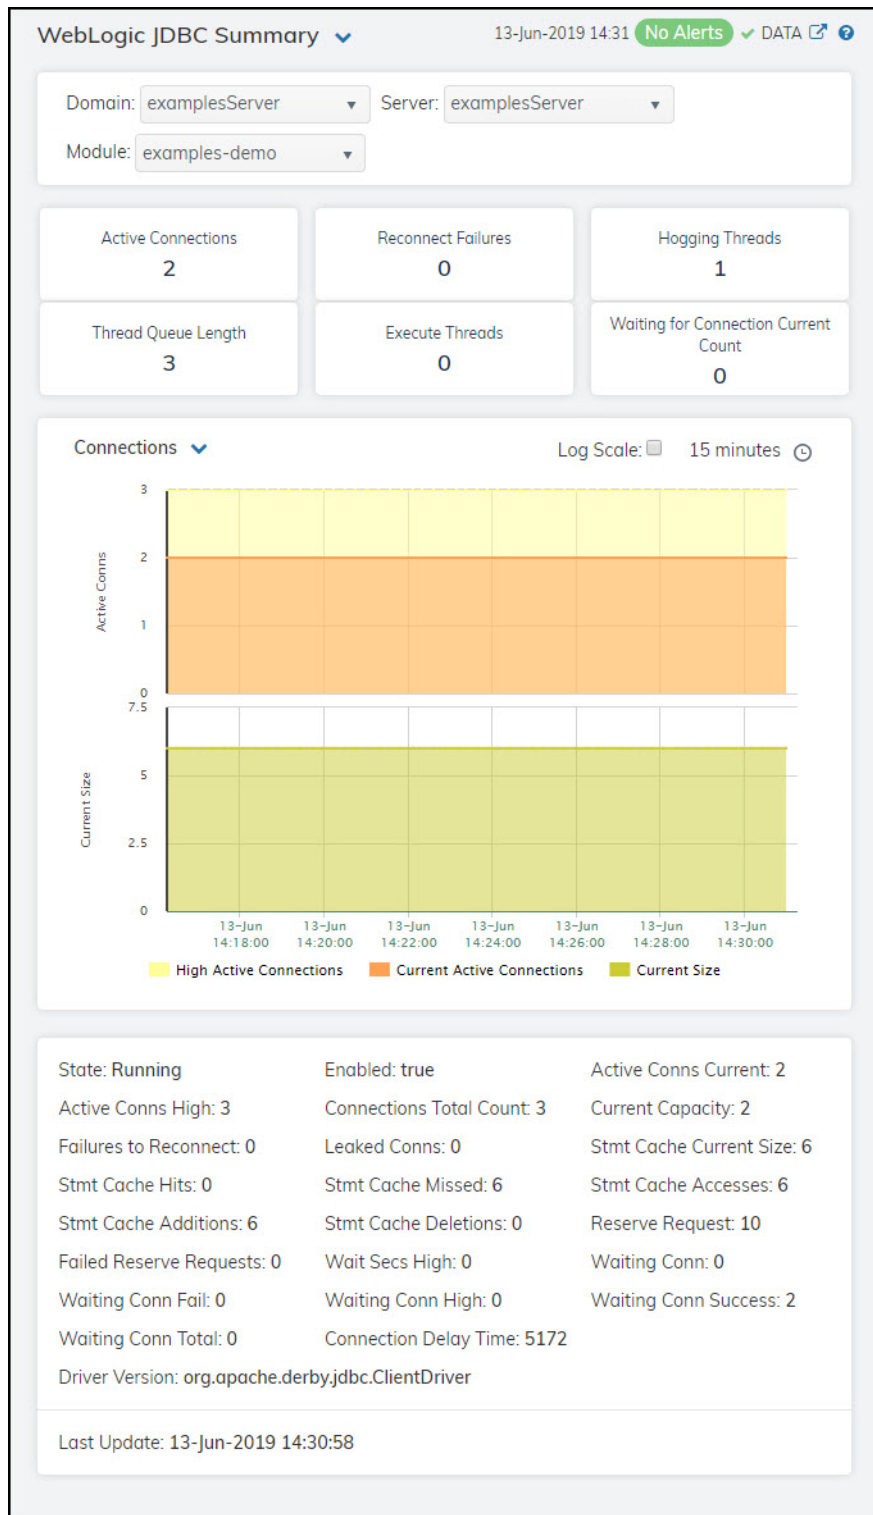
Open Sockets	The current number of sockets registered for socket muxing on this server.*
OracleHome	The Oracle home directory path.*
Parent	The name of the parent of the Java Virtual Machine.*
PendingRestartSystem Resources	The number of system resources that have not been restarted since the last restart of the WLS Server.
Restart Required	Indicates whether the server must be restarted in order to activate configuration changes.*
Restarts	The total number of restarts for this server since the cluster was last started.*
SSL Listen Address	The address on which this server is listening for SSL connections.*
Time Stamp	The date and time of the last data update.
ThreadPool Runtime Table	
Connection	The name of the connection.
Server	The name of the WLS Server located on the specified connection.
Completed Requests	The number of completed requests in the priority queue.*
Current Completed Request	The increase in the amount of completed requests (from the previous polling period to the current polling period).
Exec Idle Threads	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.*
Exec Threads Total	The total number of threads in the pool.*
Hogging Threads	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.*
Current Hogging Threads	The increase in the amount of hogging threads (from the previous polling period to the current polling period).
JRockit Runtime Table	
Connection	The name of the connection.
Server	
All Proc Avg Load	The average load of all processors in the host computer.*
Concurrent	Indicates whether the virtual machine's garbage collector runs in a separate Java thread concurrently with other Java threads.*

Free Heap	The amount, in bytes, of Java heap memory that is currently free in the virtual machine.*
Free Physical Memory	The amount, in bytes, of physical memory that is currently free on the host computer.*
GC Handles Compaction	Indicates whether the virtual machine's garbage collector compacts the Java heap.*
GcAlgorithm	The type of garbage collector (GC) that the virtual machine is using.*
Heap Free Current	The current amount of memory, in bytes, that is available in the JVM heap.*
Heap Free %	Percentage of the maximum memory that is free.*
Heap Size	The current size, in bytes, of the JVM heap.*
Max Heap Size	The maximum free memory configured for this JVM.*
Incremental	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.*
JVM Description	The description of the Java Virtual Machine.*
Java VM Vendor	The vendor of the Java Virtual Machine that the server is running.*
Java Version	The Java version of the Java Virtual Machine.*
JVM Proc Load	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.*
Jvm Type	The Java Virtual Machine type.*
time_stamp	The date and time of the last data update.
Server Version Info Region	Lists the WebLogic server version number and date that it was installed.

WebLogic JDBC Summary

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](#)" 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.



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.

Filter By:

Domain	Select the domain for which you want to view data.
Server	Select the server for which you want to view data.
Module	Select the module for which you want to see data.

Fields and Data

Active Connections	The number of connections currently in use by applications. *
Reconnect Failures	The number of times that the data source attempted to refresh a database connection and failed. *
Hogging Threads	The total number of hogging threads on this server.
Thread Queue Length	The total length of the queue thread for this server.
Execute Threads	The total number of execute threads on this server.
Waiting for Connection Current Count	The total number of connections waiting.

Trend Graphs

Connections

High Active Conns -- Traces the highest number of active database connections in the instance of the data source since the data source was instantiated.

Current 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.

Rates

Accesses/s -- Traces the cumulative, running count of the number of times that the statement cache was accessed.

Misses/s -- Traces the number of times (per second) that a statement request could not be satisfied with a statement from the cache.

Additions/s -- Traces the cumulative, running count of the number of statements added to the statement cache.


Deletions/s -- 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.

Time Settings

Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

State	The current state of the data source.
Active Conns High	The highest number of active database connections in the instance of the data source since the data source was instantiated.*
Failures to Reconnect	The number of times that the data source attempted to refresh a database connection and failed.*
Stmt Cache Hits	The cumulative, running count of the number of times statements from the cache were used.*
Stmt Cache Additions	The cumulative, running count of the number of statements added to the statement cache.*

Failed Reserve Requests	The cumulative, running count of requests for a connection from this data source that could not be fulfilled.*
Waiting Conn Fail	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.*
Waiting Conn Total	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.*
Driver Version Enabled	The driver class name of the JDBC driver used to create database connections.* Indicates whether the data source is enabled or disabled. When checked, the data source is enabled.*
Connections Total Count	The cumulative total number of database connections created in this data source since the data source was deployed.*
Leaked Conns	The number of leaked connections. *
Stmt Cache Missed	The number of times that a statement request could not be satisfied with a statement from the cache.*
Stmt Cache Deletions	The cumulative, running count of statements discarded from the cache.*
Wait Secs High	The cumulative total number of database connections created in this data source since the data source was deployed.*
Waiting Conn High	The highest number of application requests concurrently waiting for a connection from this instance of the data source.*
Connection Delay Time	The average amount of time, in milliseconds, that it takes to create a physical connection to the database.*
Active Conns Current	The number of connections currently in use by applications.*
Current Capacity	The current count of JDBC connections in the connection pool in the data source.*
Stmt Cache Current Size	The number of prepared and callable statements currently cached in the statement cache.*
Stmt Cache Accesses	The cumulative, running count of the number of times that the statement cache was accessed.*
Reserve Request	The cumulative, running count of requests for a connection from this data source.*
Waiting Conn	The number of connection requests waiting for a database connection.*
Waiting Conn Success	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.*
Last Update	The date and time of the last data update.

WebLogic Server Work Manager Table

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](#)” display, which allows you to view additional details for the server.

Name	Application Name	Completed Requests	Cur Comp. Requests
default	examples-multiDataSource-demo	0	
weblogic.wsee.mdb.DispatchPolicy	asynServletEar	0	
default	esphere	0	
weblogic.logging.DomainLogBroadc		88	
default	examples-demo	0	
default	jdbcRowSetsEar	0	
weblogic.jms.ReliableWseeSAFAge		0	
default	SamplesSearchWebApp	0	

Filter By

- Domain** Select the domain for which you want to view data.
- Server** Select the WebLogic server for which you want to see data.

Server Runtime Table

- Name** The name of the work manager.
- Application Name** The name of the application with which the work manager is associated.*
- Completed Requests** The number of requests that have been completed.*
- Cur Completed Requests** The increase in the amount of completed requests (from the previous polling period to the current polling period).
- Completed Requests/s** The rate of completed requests (per second).
- Pending Requests** The number of requests waiting in the queue.*
- Stuck Threads** The number of threads that are “stuck.”*

- Expired** 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.
- time_stamp** The date and time the data in the row was last updated.

WebLogic JMS Persistent Stores Detail Tables

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

Filter By

- Domain** Select the domain for which you want to view data.

Persistent Store Runtime Table

- Name** The name of the configuration.

Server	The name of the WLS Server located on the specified connection.
Create Count	The number of create requests issued to the persistent store.*
Delete Count	The number of delete requests issued by this persistent store.*
Health State	The health state of the store.*
Object Count	The number of objects contained in the persistent store.*
Physical Write Count	The number of times the persistent store flushes its data to durable storage.*
Read Count	The number of read requests issued to this persistent store.*
Update Count	The number of update requests issued by this persistent store.*
time_stamp	The date and time the data in the row was last updated.

Persistent Store Connection Runtime Table

Name	The name of the configuration.
Server	The name of the WLS Server located on the specified connection.
Persistent Store Runtime	The name of the persistent store.
Create Count	The number of create requests issued by this connection.*
Delete Count	The number of delete requests issued by this connection.*
Object Count	The number of objects contained in this connection.*
Read Count	The number of read requests issued by this connection.*
Update Count	The number of update requests issued by this connection.*
time_stamp	The date and time the data in the row was last updated.

WebLogic Applications View

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](#)” 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](#)” 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](#)” 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](#)” 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](#)” 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](#)” 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](#)” 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](#)” display, which allows you to view performance, utilization, and trend data for a particular application component on a WebLogic Server.

WebLogic Cluster Applications Table

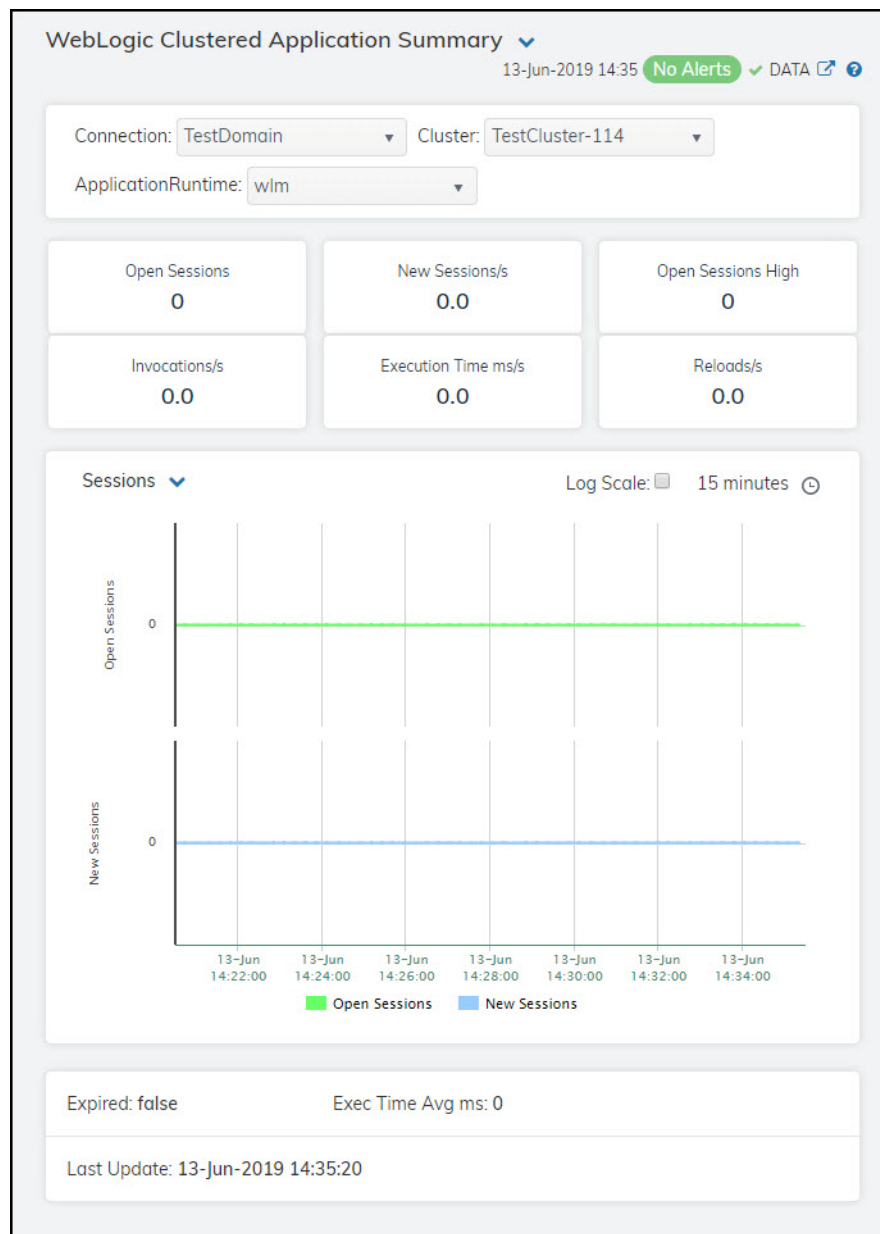
Clicking **WebLogic Applications** in the left/navigation menu opens the **WebLogic Cluster Applications Table** display, which allows you to view performance and utilization metrics for all applications on a particular cluster, or for all applications on all clusters. Double-clicking on a row in the table opens the [“WebLogic Clustered Application Summary”](#) display, which allows you to view additional information for that particular application.

Domain	Cluster	Application	Open Sessions
examplesServer		xmlBeanEar	
examplesServer		xmlBeanEar	
examplesServer		xmlBeanEar	
examplesServer		xmlBeanEar	
examplesServer		xmlBeanEar	
examplesServer		xmlBeanEar	
examplesServer		xmlBeanEar	
examplesServer		xmlBeanEar	

WebLogic Clustered Application Summary

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](#)” 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.



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

Connection	Select the connection containing the cluster for which you want to view data.
Cluster	Select the cluster for which you want to view data.
Application Runtime	Select the application for which you want to view data.

Fields and Data

Open Sessions	Displays the number of open sessions.*
New Sessions/s	The number of new sessions being created per second.
Open Sessions High	The highest number of open sessions.
Invocations/s	The number of invocations per second being created.
Execution Time ms/s	The total amount of execution time, in milliseconds per second.
Reloads/s	The number of reloads per second.

Trend Graph

Sessions

Open Sessions -- Traces the number of open sessions.

Recent New -- Traces the number of newly created sessions.

Rates

New Sessions/s -- Traces the number of sessions created per second.

Invocations/s -- Traces the number of invocations per second.

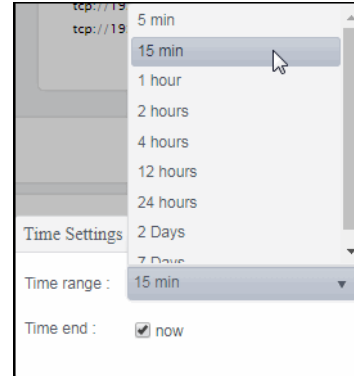
Exec Time ms/s -- Traces the execution time in milliseconds per second.

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.

Time Settings

Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

Expired

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.

Exec Time Avg ms

Displays the average execution time, in milliseconds, on the application.*

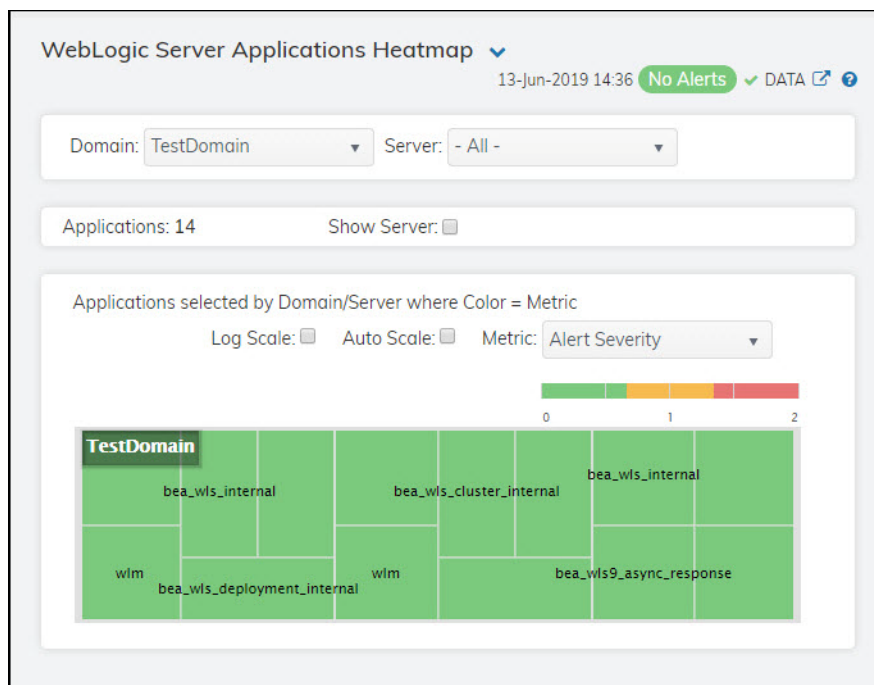
Last Update

The date and time of the last data update.

WebLogic Server Applications Heatmap

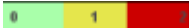

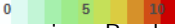

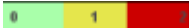

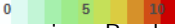

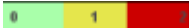

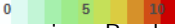

Clicking **All Apps Heatmap** in the left/navigation menu opens the **WebLogic Servers Applications Heatmap** display, which provides a heatmap view of the status and alerts of all applications within a specific WebLogic server. You can view the servers in the heatmap based on the following metrics: 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”](#) 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.



Filter By:

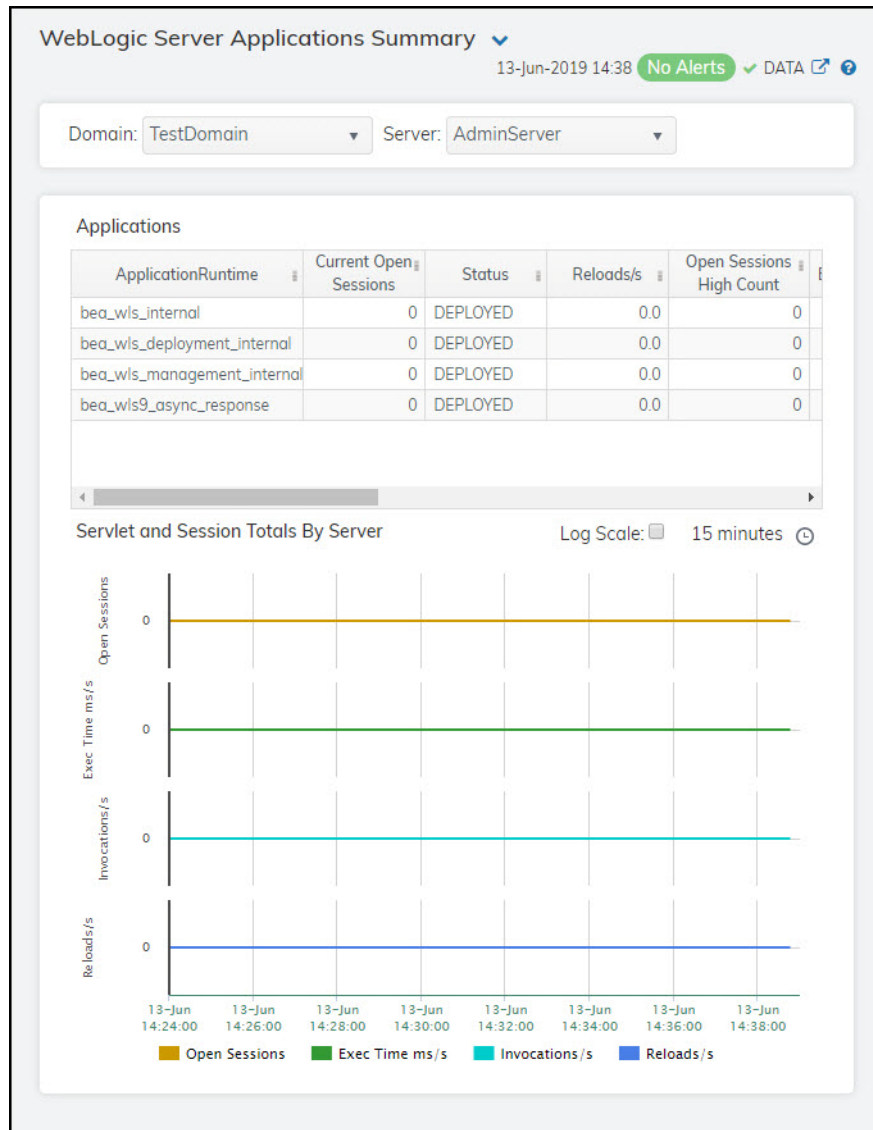
- Domain** Select the domain containing the WebLogic Server for which you want to view data.
- Server** Select the WebLogic server for which you want to view data.
- Applications** The total number of applications on the server.
- Show Server** Select this check box to display the server name in each rectangle in the heatmap.
- 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.								
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 Scale is not selected.								
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 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" display for a detailed view of metrics for that particular server. <table border="0" style="margin-top: 10px;"> <tr> <td style="vertical-align: top;">Alert Severity</td> <td>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.</td> </tr> <tr> <td style="vertical-align: top;">Alert Count</td> <td>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.</td> </tr> <tr> <td style="vertical-align: top;">Open Sessions</td> <td>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 0 to the alert threshold of WlsOpenSessionsHigh, which is 10. The middle value in the gradient bar indicates the middle value of the range (the default is 5).</td> </tr> <tr> <td style="vertical-align: top;">Open Sessions/s</td> <td>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 0 to the maximum count of objects in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</td> </tr> </table>	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.	Open Sessions	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 0 to the alert threshold of WlsOpenSessionsHigh , which is 10 . The middle value in the gradient bar indicates the middle value of the range (the default is 5).	Open Sessions/s	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 0 to the maximum count of objects in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
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.								
Open Sessions	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 0 to the alert threshold of WlsOpenSessionsHigh , which is 10 . The middle value in the gradient bar indicates the middle value of the range (the default is 5).								
Open Sessions/s	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 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/s	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.
Invocations	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 ms	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 0 to the maximum count of objects in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
Reloads	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 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 Server Applications Summary

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](#)” display, where you can view additional data for the selected application.



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.

Filter By:

- Domain** Select the domain containing the WebLogic Server for which you want to view data.
- Server** Select the WebLogic server for which you want to view data.

Applications Table

Application Runtime	The name of the application.
Current Open Sessions	The number of open sessions.
Status	The status of the application.
Reloads / sec	The rate of reloads (per second).
Open Sessions High Count	The highest number of sessions opened at one time.
Exec Time ms/s	The rate of execution time in milliseconds per second.
Invocations	The number of invocations.
Open Sessions	The number of currently open sessions.
Deployment State	The current status of the application's deployment.
Sessions Opened Total Count	The total number of opened sessions.
Invocations/s	The number of invocations per second.
Reloads	The number of reloads.
Expired	Set to true 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, this column will be set to false . This setting will remain false as long as the cache has been refreshed within the specified cache expiration time and the data is current.
time_stamp	The date and time this row of data was last updated.

Servlet and Session Totals By Server

Shows the following:


- Open Sessions** -- Traces the total number of open sessions in the application.
- Exec Time ms/s** -- Traces the execution time per second in the application.
- Invocations/s** -- Traces the number of invocations per second.
- Reloads/s** -- 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.

Time Settings

Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

WebLogic Application Summary

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”](#) 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 🕒

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


Filter By

Domain	Select the domain for which you want to view data.
Server	Select the WebLogic server for which you want to view data.
Application	Select the application for which you want to view data.
Open Sessions	The number of open sessions on the server.
Session Timeout s	The rate of session timeouts, per second.
Current Completed Requests	The increase in the amount of completed requests (from the previous polling period to the current polling period).
Pending Requests	The number of requests waiting in the queue.*
Stuck Threads	The number of threads that are "stuck."*
Exec Time Avg ms	The average execution time, in milliseconds.

Application Servlets

Servlet	The name of the servlet.
Alert Level	The current alert level.  -- 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.
Alert Count	The total number of alerts.
Completed Requests	The number of requests that have been completed.*
Cur Completed Requests	The increase in the amount of completed requests (from the previous polling period to the current polling period).
Completed Requests/s	The rate of completed requests (per second).
Pending Requests	The number of requests waiting in the queue.*
Cur Pending Requests	The increase in the amount of pending requests (from the previous polling period to the current polling period).
Pending Requests/s	The rate of pending requests (per second).
Stuck Threads	The number of threads that are "stuck."*
Cur Stuck Threads	The increase in the amount of "stuck" threads (from the previous polling period to the current polling period).

Stuck Threads/s	The rate of threads becoming “stuck” (per second).
Application Activity Trends Graph	Shows the following: <ul style="list-style-type: none"> Current Open Sessions -- Traces the total number of open sessions in the application. Exec Time ms/s -- Traces the execution time, in milliseconds per second, in the application. Invocations/s -- Traces the number of invocations per second. Reloads/s -- 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.
Time Settings	Select a time range from the drop down menu varying from 5 Minutes to Last 7 Days . By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

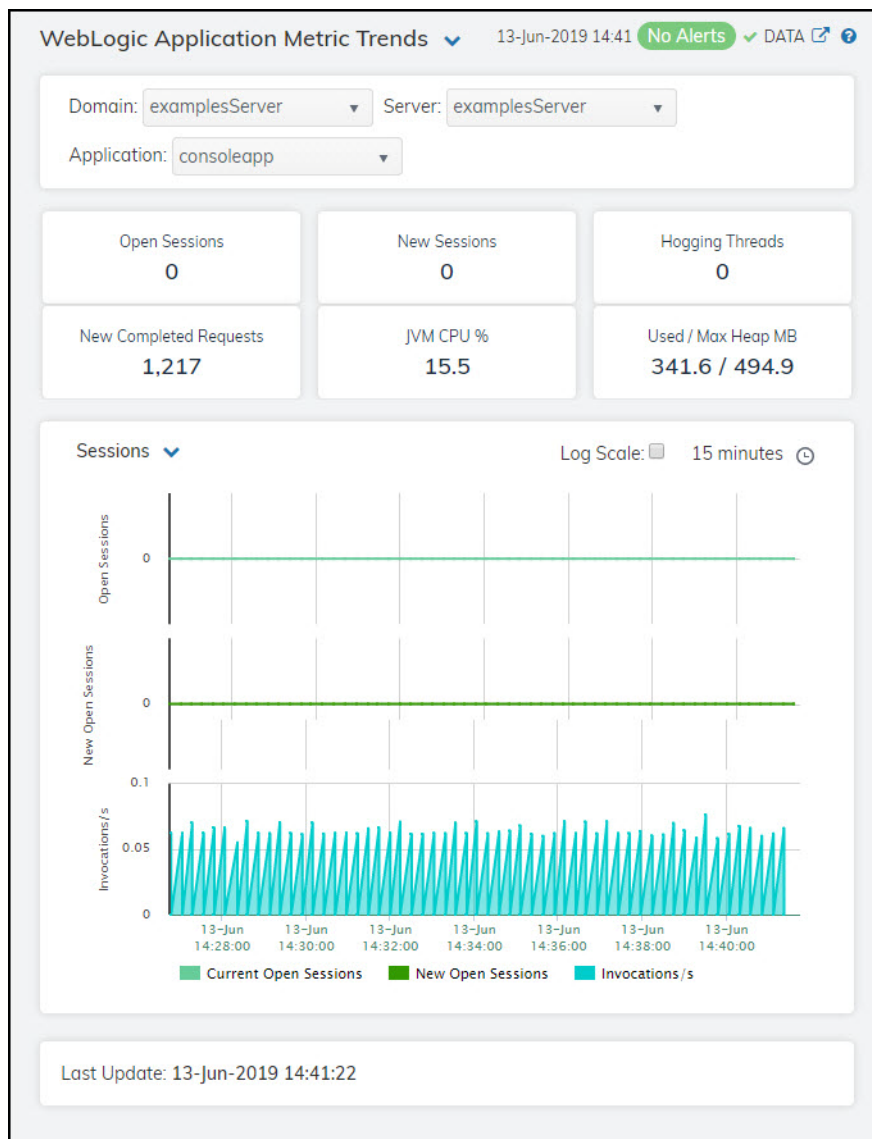
Servlets	The number of servlets.
Reloads	The number of reloads.

Open Sessions	The number of open sessions.
Exec Time Avg ms	The average execution time, in milliseconds.
Invocations	The number of invocations.
Expired	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.
Last Update	The date and time the data in the row was last updated.

WebLogic Application Metric Trends

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.



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.

Filter By:

Domain	Select the domain containing the WebLogic Server for which you want to view data.
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.

Fields and Data

Open Sessions	The total number of open sessions in the application.
New Sessions	The number of new sessions in the application.
Hogging Threads	The number of hogging threads in the application.
New Completed Requests	The number of new completed requests in the application.
JVM CPU %	The JVM CPU percentage in the application.
Used/Max Heap MB	The max heap used and the used heap, in megabytes, in the application.

Trend Graphs

Sessions

Open Sessions -- Traces the total number of open sessions in the application.

New Open Sessions -- Traces the number of new sessions in the application.

Invocations/s -- Traces the number of invocations per second in the application.

Threads

Thread Requests -- Traces the number of thread requests in the application.

Hogging Threads -- Traces the number of hogging threads in the application.

JVM

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.

Time Settings

Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

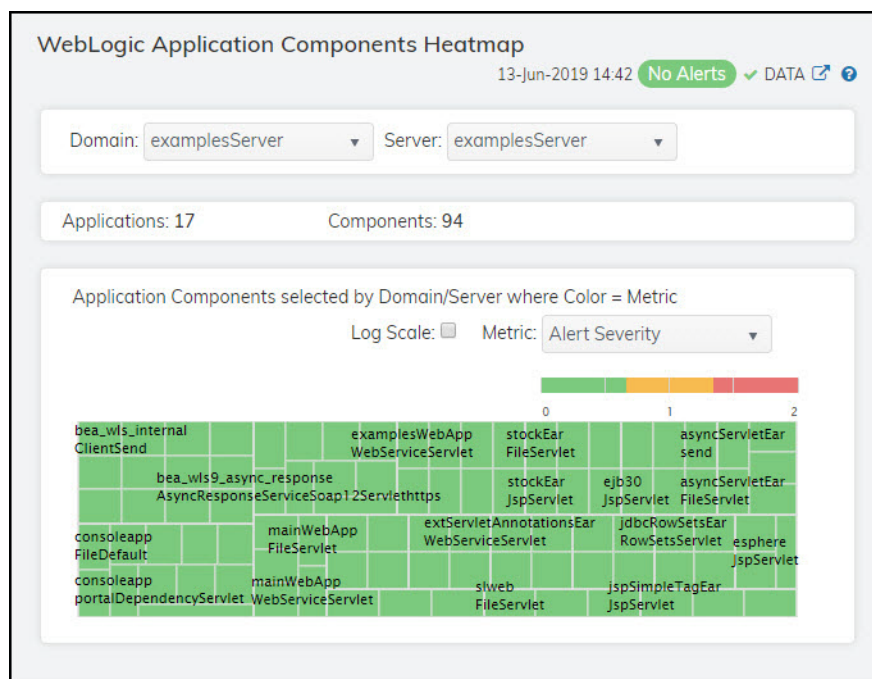
Last Update

The date and time of the last data update.

WebLogic Application Components Heatmap

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](#)” 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.



Filter By:

Domain Select the domain containing the WebLogic Server for which you want to view data.

Server Select the WebLogic server for which you want to view data.


Applications Displays the total number of application components in the WebLogic server.

Components Displays the total number of components in the WebLogic server.

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 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](#)" display for a detailed view of metrics for that particular application.


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.

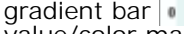
1 -- Metrics that have exceeded their specified **WARNINGLEVEL** threshold and have an Alert Severity value of **1** are shown in yellow. For a given rectangle, this indicates that one or more metrics have exceeded their warning threshold.


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.

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.

Completed Requests/s 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 **0** to the maximum count of objects in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

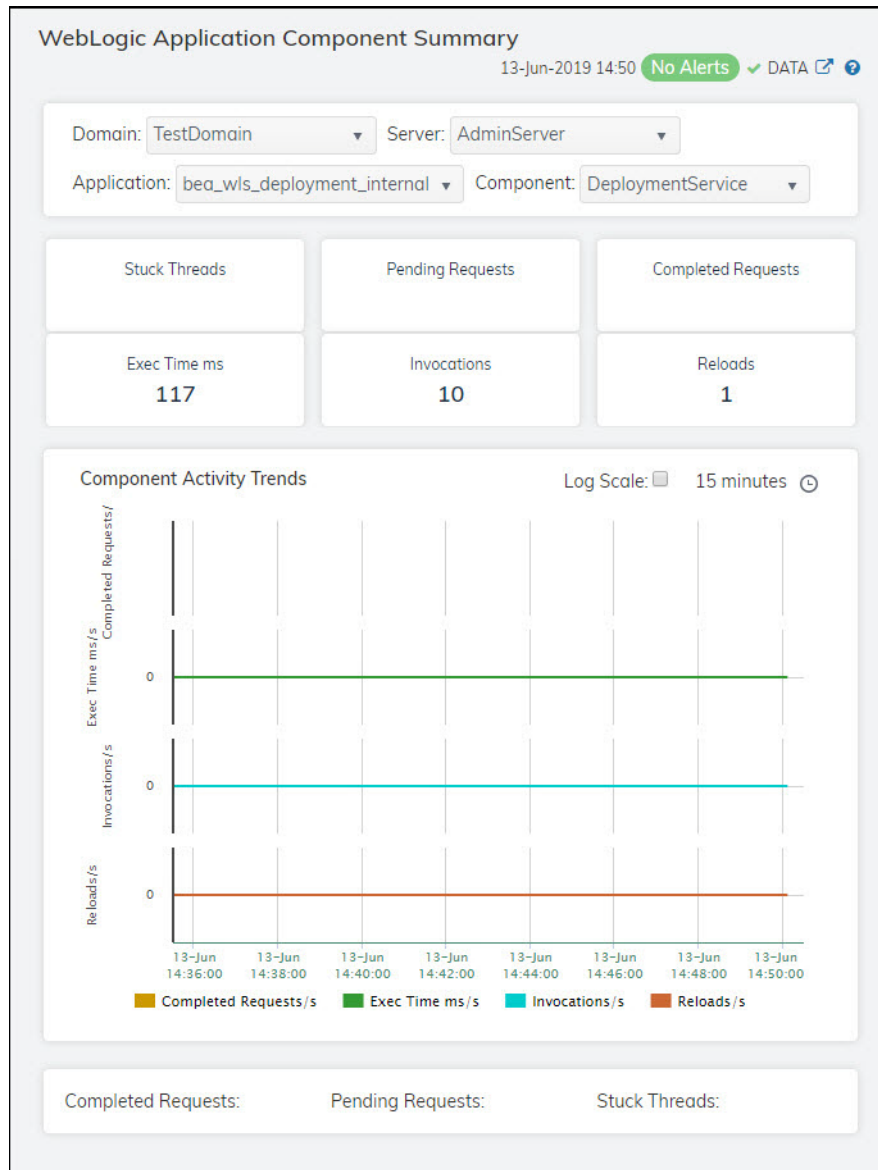
Invocations/s The number of invocations per second in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of objects in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

Invocations	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 ms	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.
Reloads	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 Application Component Summary

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.



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.

Filter By:

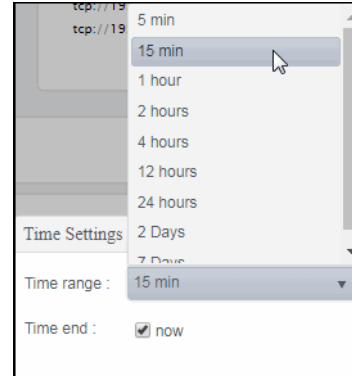
Domain	Select the domain containing the WebLogic server for which you want to view data.
Server	Select the WebLogic server containing the application for which you want to view data.
Application	Select the application for which you want to view data.
Component	Select the component for which you want to view data.

Fields and Data

Stuck Threads	The number of stuck threads.
Pending Requests	The number of pending requests.
Completed Requests	The number of completed requests.
Exec Time ms	The amount of time (in milliseconds) it took to execute the last invocation.*
Invocations	The total number of invocations of the application or component.*
Reloads	The total number of times the WebLogic server has reloaded the component since it was last deployed.*
Component Activity Trends Graph	Shows message data for the selected collection. Completed Requests/s -- Traces the rate of completed requests in the application. Exec. Time ms/s -- Traces the rate of executions, in milliseconds per second, in the application. Invocations/s -- Traces the rate of invocations in the application. Reloads/s -- Traces the rate of reloads 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.

Time Settings

Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

Completed Requests	The total number of completed requests.*
Pending Requests	The total number of pending requests.*
Stuck Threads	The total number of stuck threads.*

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 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.

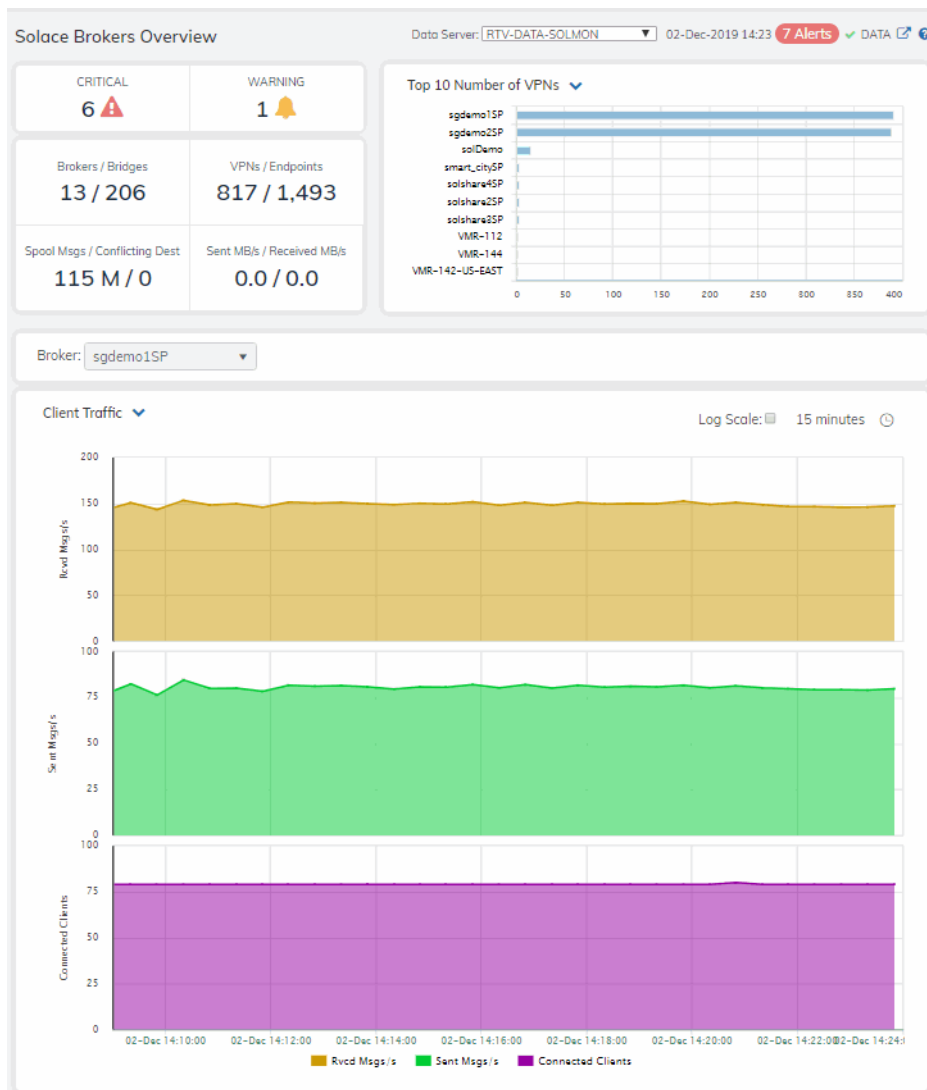
Select a data server, broker and metric from the drop-down menus. Consider keeping this display open for monitoring at a glance. You can easily view the current data for that Data Server including:

- **Top 10** most utilized **VPNs / Endpoints, Clients Connected** and **Spooled Messages**.
- The number of **Brokers / Bridges**.
- The number of **Spooled 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.
Spooled Msgs/Conflicting Dest	Total number of spooled 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

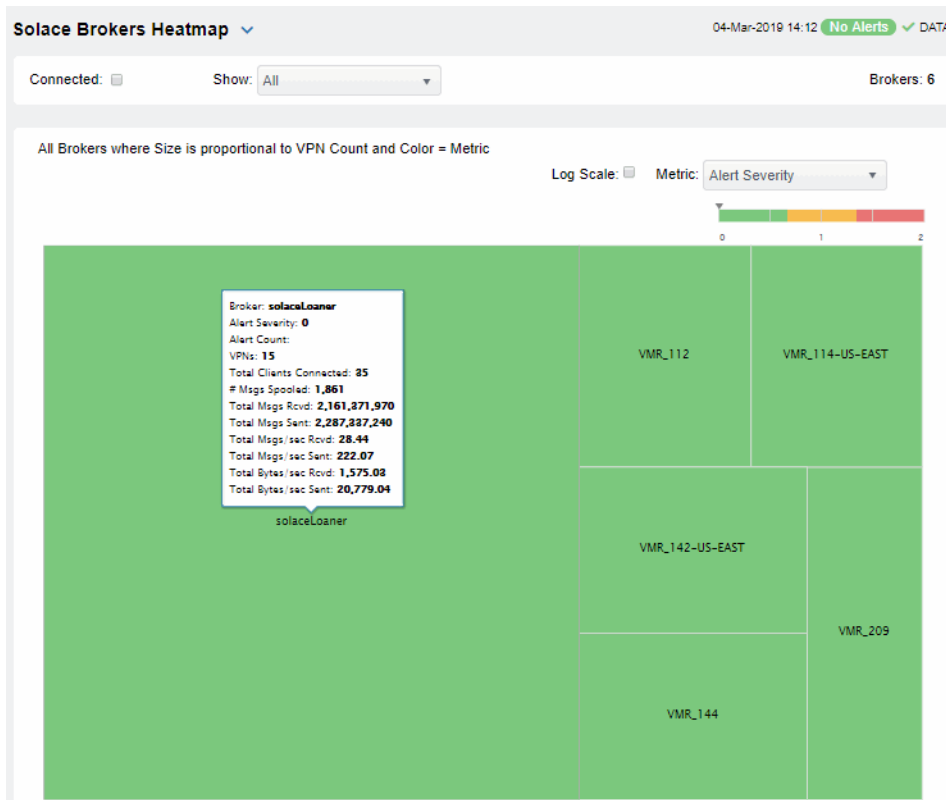
The current alert severity. Values range from **0** - **2**, as indicated in the color gradient bar, where **2** is the highest Alert Severity:













- 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.



Alert Severity	<p>The current alert severity. Values range from 0 - 2, as indicated in the color gradient  bar, where 2 is the highest Alert Severity:</p> <ul style="list-style-type: none">  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	<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 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.</p>
# Msgs Spooled	<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 0 to the defined alert threshold of SoIMsgRouterPendingMsgsHigh. The middle value in the gradient bar indicates the middle value of the range.</p>
Total Msgs Rcvd	<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 0 to the maximum count of total messages received in the heatmap. The middle value in the gradient bar indicates the average count.</p>
Total Msgs Sent	<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 0 to the maximum count of total messages sent in the heatmap. The middle value in the gradient bar indicates the average count.</p>
Total Msgs/ sec Rcvd	<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 SoIMsgRouterInboundMsgRateHigh. The middle value in the gradient bar indicates the middle value of the range.</p>
Total Msgs/ sec Sent	<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 0 to the defined alert threshold of SoIMsgRouterOutboundMsgRateHigh. The middle value in the gradient bar indicates the middle value of the range.</p>
Total Bytes/ sec Rcvd	<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 SoIMsgRouterInboundByteRateHigh. The middle value in the gradient bar indicates the middle value of the range.</p>
Total Bytes/ sec Sent	<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 SoIMsgRouterOutboundByteRateHigh. The middle value in the gradient bar indicates the middle value of the range.</p>

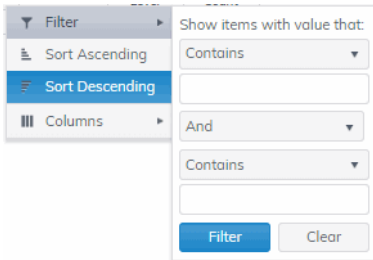
Brokers Table

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.sov		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.messag		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.mess		0	0	
VMR-112			0	ip-172-30-1-112	172.30.1.112	2	11	

Column Values

Broker	The name of the broker.
Connected	The broker state:  Red indicates that the broker is NOT connected.  Green indicates that the broker is connected.
Alert Severity	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.
Expired	When checked, performance data about the sensor has not been received within the time specified.
Host Name	The name of the host.
Platform	The name of the platform.
OS Version	The version of the operating system.
Up Time	The amount of time that the broker has been up and running.
VPNs	The total number of VPNs configured on the broker.
Total Clients	The total number of clients associated with the broker.
Total Clients Connected	The total number of clients that are currently connected to the broker.
Clients Using Compression	The number of clients who send/receive compressed messages.
Clients Using SSL	The number of clients using SSL for encrypted communications.
Max Client Connections	The maximum number of available client connections.
Endpoints	The total number of endpoints configured on the broker.
Bridges	The total number of bridges configured on the broker.
Local Bridges	The total number of local bridges configured on the broker.
Remote Bridges	The total number of remote bridges configured on the broker.
Remote Bridge Subscriptions	The total number of remote bridge subscriptions configured on the broker.
Routing Enabled	This check box is checked when the broker is configured to route messages to other brokers.
Routing Interface	The name of the interface configured to support message routing.

Total # Conflicting Destinations	The total number conflicting destinations.
SpooledSpooledMessages	The number of spooled messages on the broker.
Total Client Msgs Rcvd	The total number of client messages received on the broker.
Total Client Msgs Sent	The total number of client messages sent by the broker.
Total Client Msgs Rcvd/sec	The total number of client messages received per second by the broker.
Total Client Msgs Sent/ sec	The total number of client messages sent by the broker.
Total Client Bytes Rcvd	The total number of client bytes received by the broker.
Total Client Bytes Sent	The total number of client bytes sent by the broker.
Total Client Bytes Rcvd/sec	The total number of client bytes received per second by the broker.
Total Client Bytes Sent/sec	The total number of client bytes sent per second by the broker.
Total Client Direct Msgs Rcvd	The total number of direct client messages received by the broker.
Total Client Direct Msgs Sent	The total number of direct client messages sent from the broker.
Total Client Direct Msgs Rcvd/sec	The total number of direct client messages received per second by the broker.
Total Client Direct Msgs Sent/sec	The total number of direct client messages sent per second by the broker.
Total Client Direct Bytes Rcvd	The total number of direct client bytes received by the broker.
Total Client Direct Bytes Sent	The total number of direct client bytes sent by the broker.
Total Client Direct Bytes Rcvd/sec	The total number of direct client bytes received per second by the broker.
Total Client Direct Bytes Sent/sec	The total number of direct client bytes sent per second by the broker.
Total Client Non-Persistent Msgs Rcvd	The total number of non-persistent client messages received by the broker.
Total Client Non-Persistent Msgs Sent	The total number of non-persistent client messages sent by the broker.
Total Client Non-Persistent Msgs Rcvd/sec	The total number of non-persistent client messages received per second by the broker.
Total Client Non-Persistent Msgs Sent/ sec	The total number of non-persistent client messages sent per second by the broker.
Total Client Non-Persistent Bytes Rcvd	The total number of non-persistent client bytes received by the broker.
Total Client Non-Persistent Bytes Sent	The total number of non-persistent client bytes sent by the broker.
Total Client Non-Persistent Bytes Rcvd/sec	The total number of non-persistent client bytes received per second by the broker.
Total Client Non-Persistent Bytes Sent/sec	The total number of non-persistent client bytes sent per second by the broker.

Total Client Persistent Msgs Rcvd	The total number of persistent client messages received by the broker.
Total Client Persistent Msgs Sent	The total number of persistent client messages sent by the broker.
Total Client Persistent Msgs Rcvd/sec	The total number of persistent client messages received per second by the broker.
Total Client Persistent Msgs Sent/ sec	The total number of persistent client messages sent per second by the broker.
Total Client Persistent Bytes Rcvd	The total number of persistent client bytes received by the broker.
Total Client Persistent Bytes Sent	The total number of persistent client bytes sent by the broker.
Total Client Persistent Bytes Rcvd/sec	The total number of persistent client bytes received per second by the broker.
Total Client Persistent Bytes Sent/ sec	The total number of persistent client bytes sent per second by the broker.
Avg Egress Bytes/min	The average number of outgoing bytes per minute.
Avg Egress Compressed Msgs/min	The average number of outgoing compressed messages per minute.
Avg Egress Msgs/min	The average number of outgoing messages per minute.
Avg Egress SSL Msgs/min	The average number of outgoing messages per minute being sent via SSL-encrypted connections.
Avg Egress Uncompressed Msgs/min	The average number of uncompressed outgoing messages per minute.
Avg Ingress Bytes/min	The average number of incoming bytes per minute.
Avg Ingress Compressed Msgs/min	The average number of compressed incoming message per minute.
Avg Ingress Msgs/min	The average number of incoming messages per minute.
Average Ingress SSL Msgs/min	The average number of incoming messages per minute being received via SSL-encrypted connections.
Avg Ingress Uncompressed Msgs/min	The average number of uncompressed messages per minute.
Current Egress Bytes/sec	The current number of outgoing bytes per second.
Current Egress Compressed Msgs/sec	The current number of outgoing compressed messages per second.
Current Egress Msgs/sec	The current number of outgoing messages per second.
Current Egress SSL Msgs/sec	The current number of outgoing messages per second sent via SSL-encrypted connections.
Current Egress Uncompressed Msgs/sec	The current number of outgoing uncompressed messages per second.
Current Ingress Bytes/sec	The current number of incoming bytes per second.
Current Ingress Compressed Msgs/sec	The current number of incoming compressed messages per second.

Current Ingress Msgs/sec	The current number of incoming messages per second.
Current Ingress SSL Msgs/sec	The current number of incoming messages per second received via SSL-encrypted connections.
Current Ingress Uncompressed Msgs/sec	The current number of incoming uncompressed messages per second.
Ingress Comp Ratio	The percentage of incoming messages that are compressed.
Egress Comp Ratio	The percentage of outgoing messages that are compressed.
Egress Compressed Bytes	The number of outgoing compressed bytes.
Egress SSL Bytes	The number of outgoing compressed bytes being sent via SSL-encrypted connections.
Egress Uncompressed Bytes	The number of outgoing uncompressed bytes.
Ingress Compressed Bytes	The number of incoming compressed bytes.
Ingress SSL Bytes	The number of incoming bytes via SSL-encrypted connections.
Ingress Uncompressed Bytes	The number of incoming uncompressed bytes.
Total Egress Discards	The total number of outgoing messages that have been discarded by the broker.
Egress Discarded Msgs/sec	The total number of outgoing messages per second that have been discarded by the broker.
Total Ingress Discards	The total number of incoming messages that have been discarded by the broker.
Total Ingress Discards/sec	The total number of incoming messages per second that have been discarded by the broker.
Client Authorization Failures	The number of failed authorization attempts
Client Connect Failures (ACL)	The number of client connection failures caused because the client was not included in the defined access list.
Subscribe Topic Failures	The number of failed attempts at subscribing to topics.
TCP Fast Retrans Sent	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).
Memory (KB)	The total available memory (in kilobytes) on the broker.
Memory Free (KB)	The total amount of available memory (in kilobytes) on the broker.
Memory Used (KB)	The total amount of memory used (in kilobytes) on the broker.
Memory Used %	The percentage of total available memory that is currently being used.
Swap (KB)	The total available swap (in kilobytes) on the broker.

Swap Free (KB)	The total amount of available swap (in kilobytes) on the broker.
Swap Used (KB)	The total amount of swap used (in kilobytes) on the broker.
Swap Used %	The percentage of total available swap that is currently being used.
Subscription Mem Total (KB)	The total amount of available memory (in kilobytes) that can be used by queue/topic subscriptions.
Subscription Mem Free (KB)	The current amount of available memory (in kilobytes) that can be used by queue/topic subscriptions.
Subscription Mem Used (KB)	The current amount of memory (in kilobytes) being used by queue/topic subscriptions.
Subscription Mem Used %	The percentage of available memory being used by queue/topic subscriptions.
Chassis Product Number	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.
BIOS Version	The basic input/output system used by 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.
Operational Power Supplies	The number of available power supplies that are operational on the chassis.
Power Redundancy Config	The configuration used by the backup broker.
Max # Bridges	The maximum number of bridges allowed on the broker.
Max # Local Bridges	The maximum number of local bridges allowed on the broker.
Max # Remote Bridges	The maximum number of remote bridges allowed on the broker.
Max # Remote Bridge Subscriptions	The maximum number of remote bridge subscriptions allowed on the broker.
Redundancy Config Status	The status of the redundancy configuration.
Redundancy Status	The status of the redundant broker.
Redundancy Mode	Refer to Solace documentation for more information.
Auto-revert	Refer to Solace documentation for more information.
Mate Router Name	If redundancy is configured, this field lists the redundant broker name (mate broker name).
ADB Link Up	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.

ADB Hello Up	Refer to Solace documentation for more information.
Pair Primary Status	The primary status of the broker and its redundant (failover) mate.
Pair Backup Status	Refer to Solace documentation for more information.
CPU Cores	This value is retrieved by the "show system" SEMP request. Refer to Solace documentation for more information.
Expired	When checked, performance data about the broker has not been received within the time specified.
Time Stamp	The date and time the row of data was last updated.

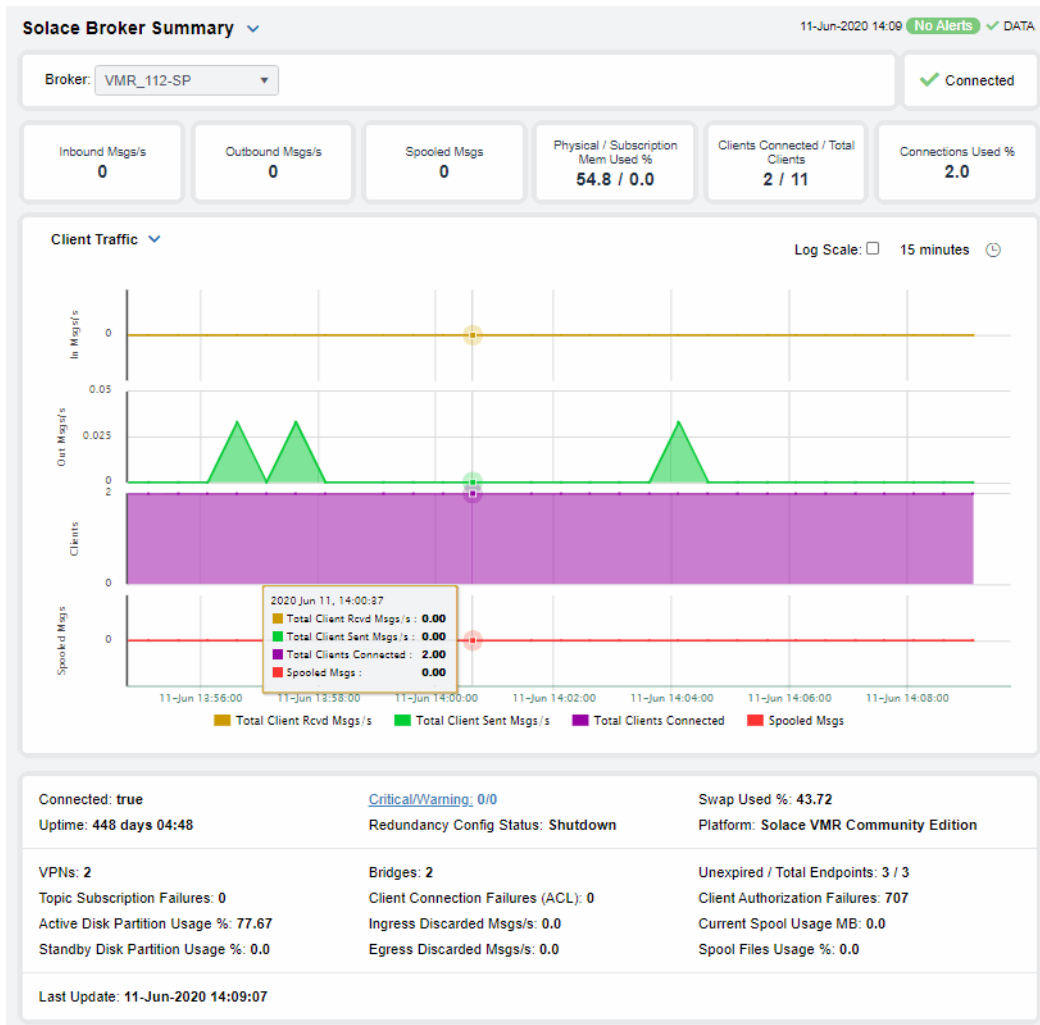
Broker Summary

View performance and processing details for a single broker, such as the total **Inbound / Outbound Messages per second**, **Spooled Messages** and **Clients Connected / Total Clients**. The trend graph traces the performance metric you select: **Client Traffic**, **Spool Msgs** or **Memory**.

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 display 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, such as connection status and **Unexpired / Total Endpoints**.

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.

Spooled Msgs/s

The number of spooled 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 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.
- **Spooled Msgs** - Traces the total number of spooled messages.

Spool Msgs

- **Spooled Msgs**- Traces the total number of spooled 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. This display only applies to Hardware (HW) Brokers. Note that the drop down menu does not show connection strings to PubSub+ Software Brokers.

Select a HW broker from the drop-down menu. 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:

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	-87.000	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

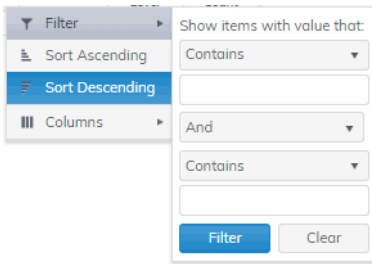
Each row in the table is a different sensor on the broker.

Type	See vendor documentation for details.
Sensor Name	The name of the sensor.
Value	Lists the value of the sensor.
Units	Lists the unit of measure for the sensor.
Status	The current status of the sensor.
Expired	When checked, performance data about the broker has not been received within the time specified.
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, operating system version, 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: VMR_112-SP

Host Name: ip-172-30-1-112	CPU-1: Not Specified	CPU-2: Not Specified
Platform: Solace VMR Community Edition	OS Version: soltr_7.2.1.617	Chassis Product Number: Xen - HVM domU
BIOS Version: 4.2.amazon	Chassis Revision: Unknown	Chassis Serial: Unknown
Total Memory (KB): 4,038,412	Memory Used (KB): 3,922,688	Memory Used %: 54.76
Swap (KB): 2,047,996	Swap Used (KB): 895,468	Swap Used %: 43.72
Operational Power Supplies: 0	Mate Router Name:	Power Redundancy Config: N/A
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-001ET-01-A	6.2.0.406	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	LFC0848B99409
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.

Mate Router Name	If redundancy is configured, this field lists the redundant broker name (mate broker name).
Configuration Status	The status of the configuration for the backup broker.
Redundancy Status	The status of the redundant broker.
Redundancy Mode	Refer to Solace documentation for more information.
Primary Status	The status of the primary broker.
Backup Status	Refer to Solace documentation for more information.
Auto-Revert	Refer to Solace documentation for more information.
ADB Link Up	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.
ADB Hello Up	Refer to Solace documentation for more information.

Fabric

Slot	Displays the slot number on the network switch.
Card Type	The type of card connected to the particular slot.
Product	The product associated with the particular slot.
Serial #	The serial number of the product.
Fw-Version	The firmware version of the product.


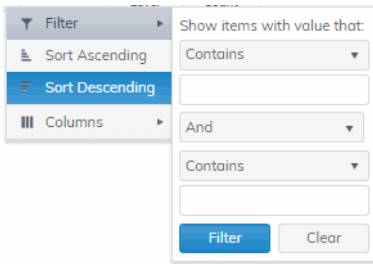
Broker Interface

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  .

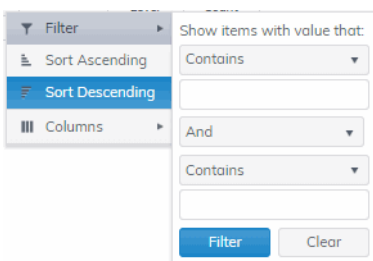


Toggle forward/backward in the trend graph per the period you choose (from the **Time range** drop-down menu) using arrows   .
Restore settings to current time by selecting **now**  .

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.

Solace Broker Spool Table 04-Mar-2019 15:35 No Alerts DATA

Broker: - All -

Count: 1

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	Topic Subscriptions on Queue Used	Max Topic Subscriptions on Queue	Sequenced Topics Used	Max Sequenced Topics	Spool Files Used	Spool Files Available
0.0	13	1	0.0	0.0	18							

Count	The number of brokers that are using spooling in the table.
Connection	The connection string associated with the broker.
Config Status	The message spool configuration status.
Operational Status	The operational status of the message spool.
Current Spool Usage (MB)	The current amount of spool used in megabytes on the broker (calculated by summing spool used for each endpoint).
Msg Spool Used By Queue	The amount of spool used by queue.
Msg Spool Used By DTE	The amount of spool used by DTE.
Message Count % Utilization	The percentage messages that use the message spool.
Delivered UnAcked Msgs % Utilization	The percentage of unacknowledged messages delivered from the message spool.
Ingress Flow Count	The current incoming flow count.
Ingress Flows Allowed	The number of incoming flows allowed.
Topic Subscriptions on Queue Used	The number of queue/topic subscriptions used.
Max Topic Subscriptions on Queue	The maximum number of queue/topic subscriptions available.
Sequenced Topics Used	The number of sequenced topics used.
Max Sequenced Topics	The maximum number of sequenced topics available.
Spool Files Used	The number of spool files used.
Spool Files Available	The maximum number of spool files available.

Spool Files % Utilization	The percentage of available spool files that have been used.
Active Disk Partition % Usage	The percentage of active disk partition that has been used.
Standby Disk Partition % Usage	The percentage of standby disk partition that has been used.
Disk Usage Current (MB)	The current amount of spool disk usage in megabytes.
Disk Usage Max (MB)	The maximum amount of spool disk usage in megabytes.
Transacted Sessions Used	The current number of transacted sessions.
Transacted Sessions Max	The maximum number of transacted sessions .
Transacted Session Count % Utilization	The percentage of transacted sessions that have been used.
Transacted Session Resource % Utilization	The percentage of transacted session resources that have been used.
Expired	When checked, performance data about the broker has not been received within the time specified.

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.
- **"CSPF 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. Each row in the table is a neighbor link. 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 [↗](#) [?](#)

Broker:

[More Columns](#)

Show Ok Only: Show: 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.
- State** The current state of the neighbor link.
- Up Time** The amount of time the broker has been up and running.
- Connections** The number of connections.
- Link Cost Actual** Refer to Solace documentation for more information.

Link Cost Configured	Refer to Solace documentation for more information.
Data Port	Refer to Solace documentation for more information.
Expired	When checked, performance data about the broker has not been received within the time specified.
Timestamp	The date and time the row of data was last updated.

CSPF Neighbors Diagram

Use this topology display to monitor the health of network components: Solace brokers, VMRs, servers and neighbor links. Quickly identify broker neighbors that are inactive and which resources their performance impacts.

Each node in the display is a Solace broker, VMR or server. If the value of the node's **Connected** column is **true**, then the node is green. If the value in its **Expired** column is **true**, it is dark gray. Otherwise the node is not connected and not expired, and the color is light gray.

No Alerts shows the number of current alerts associated with the objects and uses the following color code:

- Red indicates that there are one or more alerts in a critical state.
- Yellow indicates that there are one or more one or more alerts in a warning state.
- Green indicates that there are no alerts.

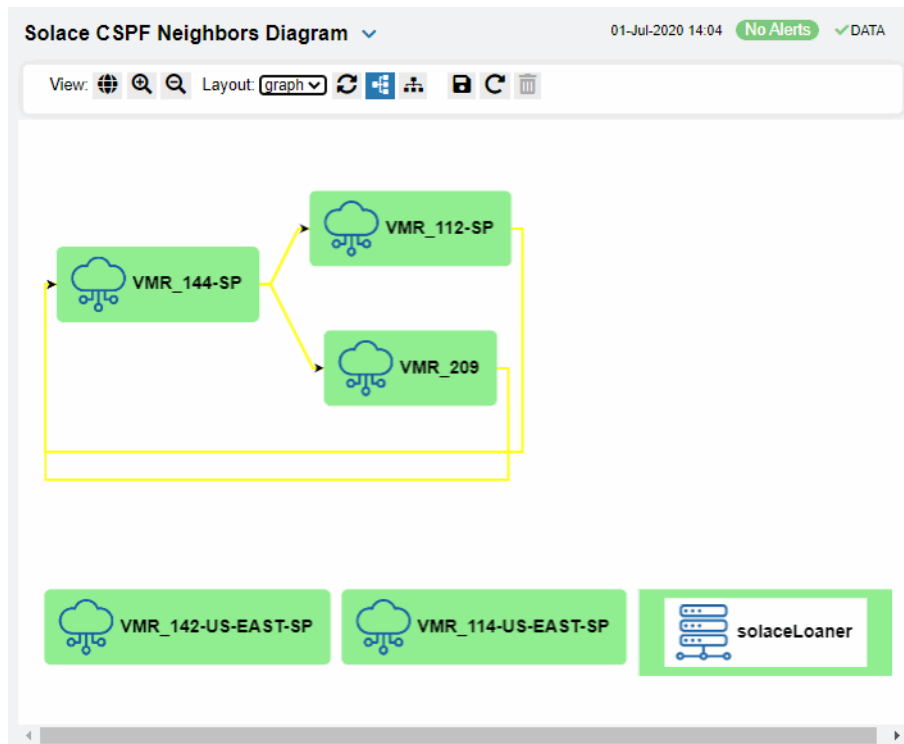
The lines connecting nodes represent neighbor links. If the value of the neighbor link **State**, is **OK** (which is shown in the "[Neighbors Table](#)"), then the neighbor link is green and if not the neighbor link yellow.

Drag and drop objects to arrange them on the screen (doing so does not logically impact Solace brokers, software or servers).


This display only shows items that are categorized as **Appliance**, **Appliance-Pair**, **VPN**, **ApplianceCapacity**, **Client**, **Bridge** or **Endpoint**. And when you click **No Alerts** to open the **Alerts Table by Component** display and investigate alerts, only alerts for items in these categories are listed.

You can mouse-over objects to see their **Host IP address** and **Platform**. Right-click on VMR objects and select **Open VMR UI** to open the Solace VMR login web page. Double-click to drill down and investigate in the "[Neighbors Table](#)".

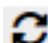
(Please scroll down for details about using the tool bar.)





Use the tool bar to interact with the diagram:


 - Toggle to include/exclude the overview (a thumbnail version of display elements that appears in the upper left margin of the display).


 - Zoom In/Zoom Out

 - Apply the selected layout options. This is useful to restore the layout after clicking and dragging the nodes or while adding links.


 - Apply a vertical flow/horizontal flow (these options apply to tree and graph layout only).

 - Save the layout options including the positions of any manually moved nodes to the local browser storage. This will be used instead of the default layout whenever you view this diagram in the browser where it was saved.

 - Restore the layout options in the toolbar to the locally saved layout if there is one, otherwise to the default layout.

 - Delete the locally saved layout options.

You can select a layout type to apply to the diagram. You can also click and drag the component nodes to lay them out manually.



The **tree** and **graph** layouts are useful on diagrams where the flow (links) between nodes go mostly in one direction, from one or more root nodes to child nodes, then to grandchild nodes, and so forth. The graph layout attempts to keep the nodes arranged neatly in rows or columns, the tree layout places the nodes more closely together.

The **net** layout can be useful for diagrams where the flow between nodes goes in multiple directions.

The **circle** layout can be useful for diagrams where the flow is mostly in one direction but where there are also links from the deepest child nodes back to top level nodes.

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.



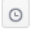
Egress Discards/s The total number of discarded messages sent from the selected **Broker** to the selected **Neighbor** broker since the broker was last started.


Trend Graphs



Traces the rates of messages sent from the selected **Broker** to the selected **Neighbor** broker.

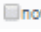
Sent Msgs/s Refer to Solace documentation for more information.

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:

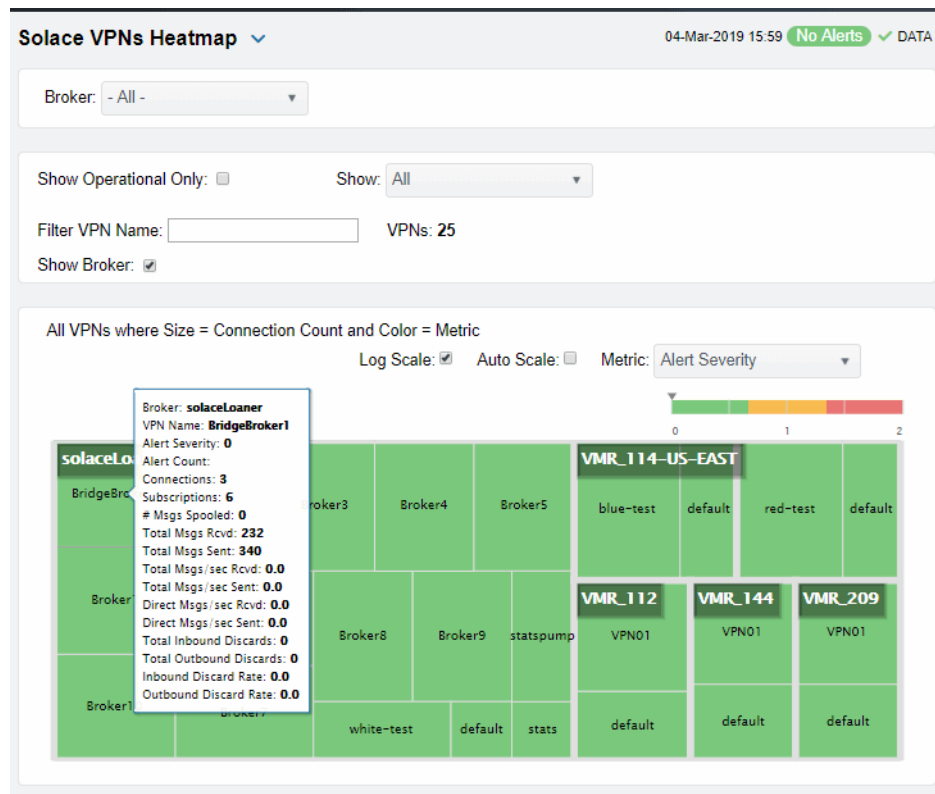
- [“VPNs Heatmap” on page 450](#): A color-coded heatmap view of the current status of all VPNs configured on a specific broker.
- [“VPNs Table” on page 454](#): A tabular view of all available data for all VPNs configured on a specific broker.
- [“VPNs Summary” on page 457](#): Current and historical 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.




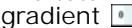


Alert Severity

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	<p>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.</p> <p>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.</p>
Subscriptions	<p>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.</p> <p>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.</p>
# Msgs Spooled	<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 0 to the defined alert threshold of SolMsgRouterPendingMsgsHigh. The middle value in the gradient bar indicates the middle value of the range.</p> <p>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.</p>
Total Msgs Rcvd	<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 0 to the maximum count of messages received in the heatmap. The middle value in the gradient bar indicates the average count.</p> <p>The Auto flag does not impact this metric.</p>
Total Msgs Sent	<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 0 to the maximum count of messages sent in the heatmap. The middle value in the gradient bar indicates the average count.</p> <p>The Auto flag does not impact this metric.</p>
Total Msgs/ sec Rcvd	<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 SolVpnInboundMsgRateHigh. The middle value in the gradient bar indicates the middle value of the range.</p> <p>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.</p>

Total Msgs/ sec Sent	<p>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.</p>
Total Bytes/ sec Rcvd	<p>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.</p>
Total Bytes/ sec Sent	<p>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.</p>
Direct Msgs/sec Rcvd	<p>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.</p>
Direct Msgs/sec Sent	<p>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.</p>
Total Inbound Discards	<p>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.</p>

Total Outbound Discards	<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 0 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 Auto flag does not impact this metric.</p>
Inbound Discard Rate	<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 SolVpnInboundDiscardRateHigh. The middle value in the gradient bar indicates the middle value of the range.</p> <p>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.</p>
Outbound Discard Rate	<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 SolVpnOutboundDiscardRateHigh. The middle value in the gradient bar indicates the middle value of the range.</p> <p>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.</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.

Operational	When checked, this status indicates that the VPN is enabled and is operating normally.
Total Unique Subscriptions	The total number of unique subscriptions to the VPN.
Total Client Messages Rcvd	The total number of messages received from clients connected to the VPN.
Total Client Messages Sent	The total number of messages sent to clients connected to the VPN.
Total Client Bytes Rcvd	The total number of bytes contained in messages received from clients connected to the VPN.
Total Client Bytes Sent	The total number of bytes contained in messages sent to clients connected to the VPN.
Total Client Msgs/sec Rcvd	The total number of messages received per second from clients connected to the VPN.
Total Client Msgs /sec Sent	The total number of messages sent per second to clients connected to the VPN.
Total Client Bytes/sec Rcvd	The total number of bytes contained in messages received per second from clients connected to the VPN.
Total Client Bytes/sec Sent	The total number of bytes contained in messages sent per second to clients connected to the VPN.
Client Direct Msgs Rcvd	The total number of direct messages received from clients connected to the VPN.
Client Direct Msgs Sent	The total number of direct messages sent to clients connected to the VPN.
Client Direct Bytes Rcvd	The total number of bytes contained in direct messages received from clients connected to the VPN.
Client Direct Bytes Sent	The total number of bytes contained in direct messages sent to clients connected to the VPN.
Client Direct Msgs/sec Rcvd	The total number of direct messages received per second from clients connected to the VPN.
Client Direct Msgs/sec Sent	The total number of direct messages sent per second to clients connected to the VPN.
Client Direct Bytes/sec Rcvd	The total number of bytes contained in the direct messages received per second from clients connected to the VPN.
Client Direct Bytes/sec Sent	The total number of bytes contained in the direct messages sent per second to clients connected to the VPN.
Client NonPersistent Msgs Rcvd	The total number of non-persistent messages received from clients connected to the VPN.
Client NonPersistent Msgs Sent	The total number of non-persistent messages sent to clients connected to the VPN.
Client NonPersistent Bytes Rcvd	The total number of bytes contained in the non-persistent messages received from clients connected to the VPN.
Client NonPersistent Bytes Sent	The total number of bytes contained in the non-persistent messages sent per second to clients connected to the VPN.
Client NonPersistent Msgs/sec Rcvd	The total number of non-persistent messages received per second from clients connected to the VPN.
Client NonPersistent Msgs/sec Sent	The total number of non-persistent messages sent per second to clients connected to the VPN.

Client NonPersistent Bytes/sec Rcvd	The total number of bytes contained in the non-persistent messages received per second from clients connected to the VPN.
Client NonPersistent Bytes/sec Sent	The total number of bytes contained in the non-persistent messages sent per second to clients connected to the VPN.
Client Persistent Msgs Rcvd	The total number of persistent messages received from clients connected to the VPN.
Client Persistent Msgs Sent	The total number of persistent messages sent to clients connected to the VPN.
Client Persistent Bytes Rcvd	The total number of bytes contained in persistent messages received from clients connected to the VPN.
Client Persistent Bytes Sent	The total number of bytes contained in persistent messages sent to clients connected to the VPN.
Client Persistent Msgs/sec Rcvd	The total number of persistent messages received per second from clients connected to the VPN.
Client Persistent Msgs/sec Sent	The total number of persistent messages sent per second to clients connected to the VPN.
Client Persistent Bytes/sec Rcvd	The total number of bytes contained in the persistent messages received per second from clients connected to the VPN.
Client Persistent Bytes/sec Sent	The total number of bytes contained in the persistent messages sent per second to clients connected to the VPN.
Total In Discards	The total number of discarded incoming messages.
Total In Discards/sec	The number of discarded incoming messages per second.
Total Out Discards	The total number of discarded outgoing messages.
Total Out Discards/sec	The number of discarded outgoing messages per second.
Max Spool Usage (MB)	The maximum amount of disk storage (in megabytes) that can be consumed by all spooled message on the VPN.
Authentication Type	The defined authentication type on the VPN.
Expired	When checked, performance data about the broker has not been received within the time specified.
Time Stamp	The date and time the row data was last updated.

VPNs Summary

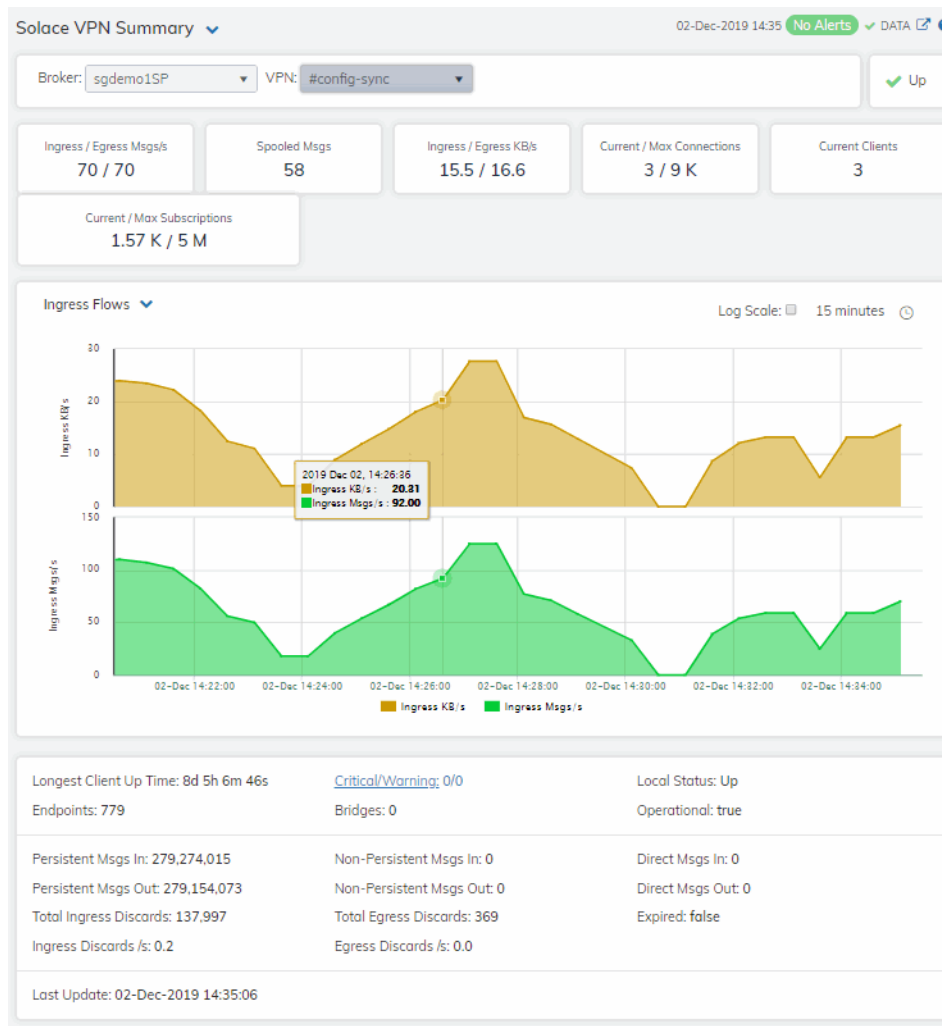
View neighbor broker current configuration details and message throughput rates.

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.
- Spoiled Msgs** The number of spoiled 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.

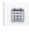

- **Spooled Msgs:** The number of spooled messages for the VPN.
- **Client Msgs/sec:** The rate of incoming messages (per second) from client.
- **Direct Client Msgs/sec:** The rate of direct incoming messages (per second) from the direct client.


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.

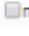
Time Settings

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.

Total In Discards/sec

The total number of incoming messages that were discarded, per second.

Critical/Warning

The number of critical alerts / warning alerts which also opens the **Alerts Table**.

Bridges

The number of bridges.

Non-Persistent Msgs In

The total number of incoming non-persistent messages.

Non-Persistent In Msgs/s	The number of incoming non-persistent messages per second.
Non-Persistent Msgs Out	The total number of outgoing non-persistent messages.
Non-Persistent Out Msgs/s	The number of outgoing non-existent messages per second.
Total Out Discards	The total number of outgoing messages that were discarded.
Total Out Discards/sec	The total number of outgoing messages that were discarded, per second.
Direct Msgs In	The total number of incoming direct messages.
Direct In Msgs/s	The number of incoming direct messages per second.
Direct Msgs Out	The total number of outgoing direct messages.
Direct Out Msgs/s	The number of outgoing direct messages per second.
Expired	When true , performance data about the VPN has not been received within the time specified.
Last Update	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	Slow Subscriber
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/v/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/8798/14	✓		

Page 1 of 3 1 - 40 of 101 items

Broker

Lists the name of the selected broker.

VPN

Lists the name of the selected VPN.

Client Name

The name of the client.

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

Total number of alerts for the client.

Slow Subscriber

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).

Total Egress Flows

The total number of outgoing flows.

Total Ingress Flows

The total number of incoming flows.

Subscriptions

The total number of subscriptions.

Subscription Msgs Rcvd

The total number of messages received from subscriptions.

Subscription Msgs Sent	The total number of messages sent from subscriptions.
Type	Lists the type of alert.
Uptime	Lists the amount of time the client has been up and running.
Client ID	Lists the client ID.
Client UserName	Lists the user name for the client.
Client Address	The IP Address of the client.
Client Profile	The client profile that is assigned to the client.
ACL Profile	The access control list profile to which the client is assigned.
Description	Lists a description of the client.
Platform	Lists the platform of the client.
Software Version	The version of the platform.
Total Flows Out	The total number of outbound message flows for the client.
Total Flows In	The total number of inbound message flows for the client.
# Subscriptions	The number of subscribers connected to the client.
Add Sub Msgs Rcvd	The number of Add Subscription messages received.
Add Sub Msgs Sent	The number of Add Subscription Messages sent.
Already Exists Msgs Sent	Refer to Solace documentation for more information.
Assured Ctrl Msgs Rcvd	Refer to Solace documentation for more information.
Assured Ctrl Msgs Sent	Refer to Solace documentation for more information.
Total Client Msgs Rcvd	The total number of messages received by the client.
Total Client Msgs Sent	The total number of messages sent by the client.
Total Client Bytes Rcvd	The total number of bytes contained within the messages received by the client.
Total Client Bytes Sent	The total number of bytes contained within the messages sent by the client.
Total Client Msgs Rcvd/sec	The total number of messages received per second by the client.
Total Client Msgs Sent/sec	The total number of messages sent per second by the client.
Total Client Bytes Rcvd/sec	The total number of bytes contained within the messages received per second by the client.
Total Client Bytes Sent/sec	The total number of bytes contained within the messages sent per second by the client.
Ctl Bytes Rcvd	The number of control data bytes received by the client.
CTL Bytes Sent	The number of control data bytes sent by the client.

Ctl Msgs Rcvd	The number of control data messages received by the client.
Ctl Msgs Sent	The number of control data messages sent by the client.
Client Data Bytes Rcvd	The number of bytes contained within the data messages received by the client.
Client Data Bytes Sent	The number of bytes contained within the data messages sent by the client.
Client Data Msgs Rcvd	The number of data messages received by the client.
Client Data Msgs Sent	The number of data messages sent by the client.
Client Direct Msgs Rcvd	The number of direct messages received by the client.
Client Direct Msgs Sent	The number of direct messages sent by the client.
Client Direct Bytes Rcvd	The number of bytes contained within direct messages received by the client.
Client Direct Bytes Sent	The number of bytes contained within direct messages sent by the client.
Client Direct Msgs Rcvd/sec	The number of direct messages received per second by the client.
Client Direct Msgs Sent/sec	The number of direct messages sent per second by the client.
Client Direct Bytes Rcvd/sec	The number of bytes contained within the messages received per second by the client.
Client Direct Bytes Sent/sec	The number of bytes contained within the messages sent per second by the client.
Client NonPersistent Msgs Rcvd	The number of non-persistent messages received by the client.
Client NonPersistent Msgs Sent	The number of non-persistent messages sent by the client.
Client NonPersistent Bytes Rcvd	The number of bytes contained within the non-persistent messages received by the client.
Client NonPersistent Bytes Sent	The number of bytes contained within the non-persistent messages sent by the client.
Client NonPersistent Msgs Rcvd/sec	The number of non-persistent messages received per second by the client.
Client NonPersistent Msgs Sent/sec	The number of non-persistent messages sent per second by the client.
Client NonPersistent Bytes Rcvd/sec	The number of bytes contained within the non-persistent messages received per second by the client
Client NonPersistent Bytes Sent/sec	The number of bytes contained within the non-persistent messages sent per second by the client
Client Persistent Msgs Rcvd	The number of persistent messages received by the client.
Client Persistent Msgs Sent	The number of persistent messages sent by the client.
Client Persistent Bytes Rcvd	The number of bytes contained within the persistent messages received by the client.
Client Persistent Bytes Sent	The number of bytes contained within the persistent messages sent by the client.
Client Persistent Msgs Rcvd/sec	The number of persistent messages received per second by the client.

Client Persistent Msgs Sent/sec	The number of persistent messages sent per second by the client.
Client Persistent Bytes Rcvd/sec	The number of bytes contained within the persistent messages received per second by the client.
Client Persistent Bytes Sent/sec	The number of bytes contained within the persistent messages sent per second by the client.
Denied Dup Clients	Refer to Solace documentation for more information.
Denied Subscribe Permission	The number of denied subscription requests due to improper permissions.
Denied Subscribe Topic-ACL	The number of denied subscriptions to topics due to the fact that the client requesting was not on the Access Control List.
Denied Unsubscribe Permission	The number of denied unsubscribe requests due to improper permissions.
Denied Unsubscribe Topic-ACL	The number of denied unsubscribe requests to topics due to the fact that the client requesting was not on the Access Control List.
DTO Msgs Rcvd	The number of Deliver-To-One messages received by the client.
Egress Compressed Bytes	The number of compressed bytes contained within outgoing messages.
Ingress Compressed Bytes	The number of compressed bytes contained within incoming messages.
Total Ingress Discards	The total number of discarded incoming messages.
Total Egress Discards	The total number of discarded outgoing messages.
Total Ingress Discards/sec	The total number of discarded incoming messages per second.
Total Egress Discards/sec	The total number of discarded outgoing messages per second.
Keepalive Msgs Rcvd	The number of Keepalive messages received by the client.
Keepalive Msgs Sent	The number of Keepalive messages sent by the client.
Large Msgs Rcvd	The number of large messages received by the client.
Login Msgs Rcvd	The number of login message received by the client.
Max Exceeded Msgs Sent	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.
Not Enough Space Msgs Sent	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.
Not Found Msgs Sent	Refer to Solace documentation for more information.
Parse Error on Add Msgs Sent	Refer to Solace documentation for more information.
Parse Error on Remove Msgs Sent	Refer to Solace documentation for more information.
Remove Subscription Msgs Rcvd	The number of remove subscription requests received by the client.

Remove Subscription Msgs Sent	The number of remove subscription requests sent by the client.
Subscribe Client Not Found	The number of subscription requests for clients that were not found.
Unsubscribe Client Not Found	The number of unsubscribe requests for clients that were not found.
Update Msgs Rcvd	Refer to Solace documentation for more information.
Update Msgs Sent	Refer to Solace documentation for more information.
Expired	When checked, performance data about the client has not been received within the time specified.
Timestamp	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.

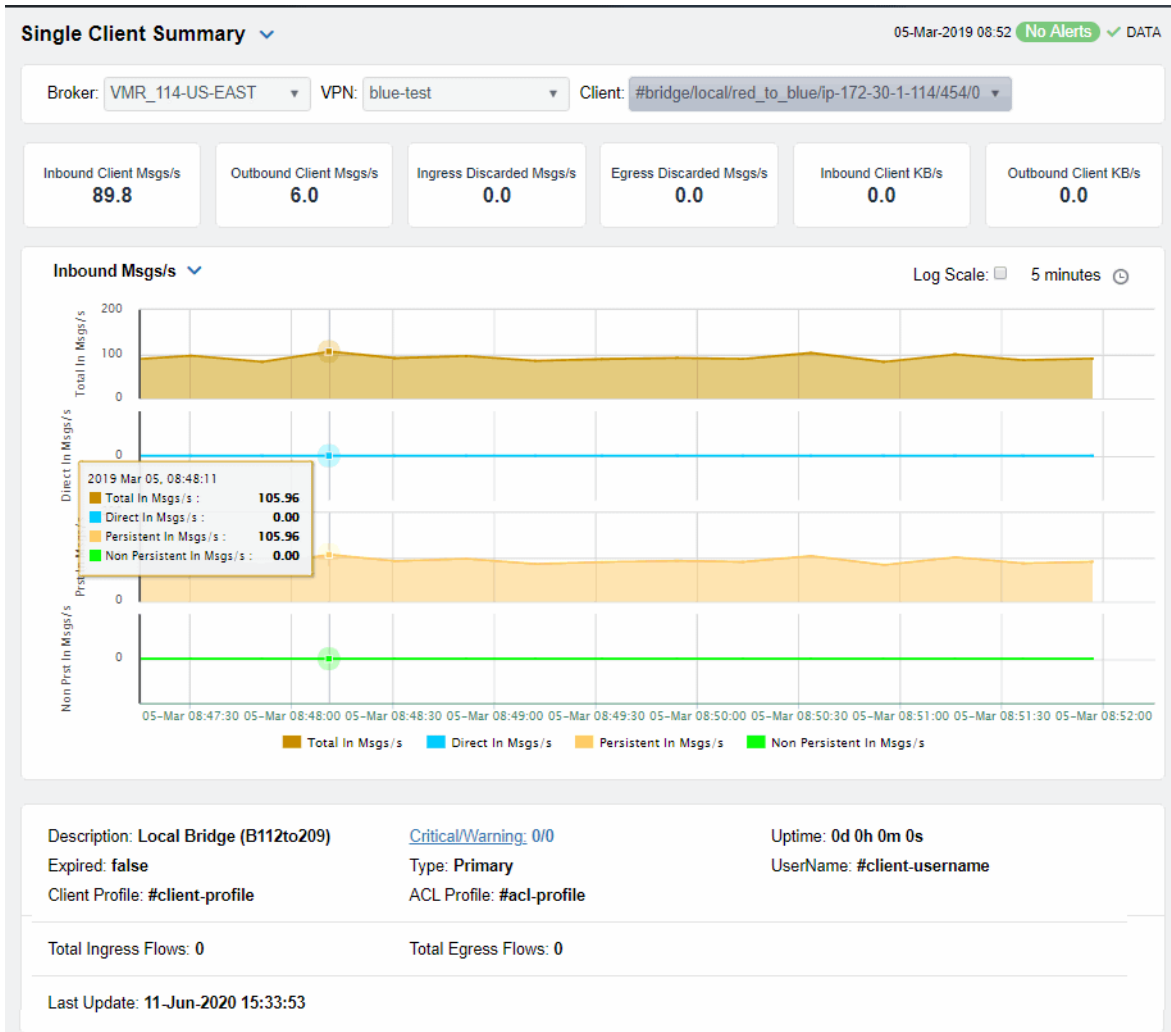
Select a broker, VPN and client from the drop-down menus. You can view the **Client Type**, the **User Name**, the **Client ID**, the associated **Platform**, the current **Up Time**, and additional information specific to the client. You can also view the total number of incoming and outgoing messages, 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 broker has not been received within the time specified.

Client Profile

The client profile that is assigned to the client.

Total Ingress Flows

The number of inflows coming to the client.

Persistent Msgs In/sec

The number of persistent incoming messages per second.

Persistent Msgs Out/sec

The number of persistent outgoing messages per second.

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**.

Non Persistent Msgs In/sec

The number of non-persistent incoming messages per second.

NonPersistent Msgs Out/sec

The number of non-persistent outgoing messages per second.

Uptime

If the VPN's **Local Status** is **Up**, this field displays the length of time that the VPN has been up and running.

Username	The client's user name.
Direct In Msgs /sec	The number of non-persistent incoming messages per second.
Direct Out Msgs /sec	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 health status and allows you to open the Solace PubSub+ Manager.
- **"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:

The screenshot shows a dialog box with the following elements:

- A left-hand menu with options: Filter (selected), Sort Ascending, Sort Descending (highlighted), and Columns.
- A main area titled "Show items with value that:" containing two search criteria sections. Each section has a dropdown menu set to "Contains" and an empty text input field.
- Buttons for "Filter" and "Clear" at the bottom.

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 Router
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.

Connection Establisher	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.
Redundancy	Displays whether the bridge is the primary bridge, the backup bridge, the static bridge (default bridge used when no other bridge is available), or whether it is the only bridge available (none).
Uptime	The current amount of time in which the bridge has been up and running.
Client Name	The name of the client.
Connected Via Addr	The local IP address and port used for the bridge.
Connected Via Interface	The name of the network interface used for the bridge.
Client Direct Bytes Rcvd	The number of bytes contained within direct messages received by the client via the bridge.
Client Direct Bytes/sec Rcvd	The number of bytes contained within direct messages received per second by the client via the bridge.
Client Direct Bytes Sent	The number of bytes contained within direct messages sent by the client via the bridge.
Client Direct Bytes/sec Sent	The number of bytes contained within direct messages sent per second by the client via the bridge.
Client Direct Msgs/sec Rcvd	The number of bytes contained within direct messages received per second by the client via the bridge.
Client Direct Msgs Sent	The number of direct messages sent by the client via the bridge.
Client Direct Msgs/sec Sent	The number of direct messages sent per second by the client via the bridge.
Client NonPersistent Bytes Rcvd	The number of bytes contained within non-persistent messages received by the client via the bridge.
Client NonPersistent Bytes/sec Rcvd	The number of bytes contained within non-persistent messages received per second by the client via the bridge.
Client NonPersistent Bytes Sent	The number of bytes contained within non-persistent messages sent by the client via the bridge.
Client NonPersistent Bytes/sec Sent	The number of bytes contained within non-persistent messages sent per second by the client via the bridge.
Client NonPersistent Msgs Rcvd	The number of non-persistent messages received by the client via the bridge.
Client NonPersistent Msgs/sec Rcvd	The number of non-persistent messages received per second by the client via the bridge.
Client NonPersistent Msgs Sent	The number of non-persistent messages sent by the client via the bridge.
Client NonPersistent Msgs/sec Sent	The number of non-persistent messages sent per second by the client via the bridge.
Client Persistent Bytes Rcvd	The number of bytes contained within persistent messages received by the client via the bridge.
Client Persistent Bytes/sec Rcvd	The number of bytes contained within persistent messages received per second by the client via the bridge.
Client Persistent Bytes Sent	The number of bytes contained within persistent messages sent by the client via the bridge.

Client Persistent Bytes/sec Sent	The number of bytes contained within persistent messages sent per second by the client via the bridge.
Client Persistent Msgs Rcvd	The number of persistent messages received by the client via the bridge.
Client Persistent Msgs /sec Rcvd	The number of persistent messages received per second by the client via the bridge.
Client Persistent Msgs Sent	The number of persistent messages sent by the client via the bridge.
Client Persistent Msgs/sec Sent	The number of persistent messages sent per second by the client via the bridge.
Total Client Bytes Rcvd	The number of bytes contained within all messages received by the client via the bridge.
Total Client Bytes/sec Rcvd	The number of bytes contained within all messages received per second by the client via the bridge.
Total Client Bytes Sent	The number of bytes contained within all messages sent by the client via the bridge.
Total Client Bytes/sec Sent	The number of bytes contained within all messages sent per second by the client via the bridge.
Total Client Msgs Rcvd	The total number of all messages received by the client via the bridge.
Total Client Msgs/sec Rcvd	The total number of all messages received per second by the client via the bridge.
Total Client Msgs Sent	The total number of all messages sent by the client via the bridge.
Total Client Msgs/sec Sent	The total number of all messages sent per second by the client via the bridge.
Total Out Discards	The total number of discarded outgoing messages sent by the client via the bridge.
Total Out Discards/sec	The total number of discarded outgoing messages sent per second by the client via the bridge.
Total In Discards	The total number of discarded incoming messages received by the client via the bridge.
Total In Discards/sec	The total number of discarded incoming messages received per second by the client via the bridge.
Expired	When checked, performance data about the broker has not been received within the time specified.
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. Open the Solace PubSub+ Manager by right-clicking on a router and selecting **Launch PubSub+ Manager**.

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.

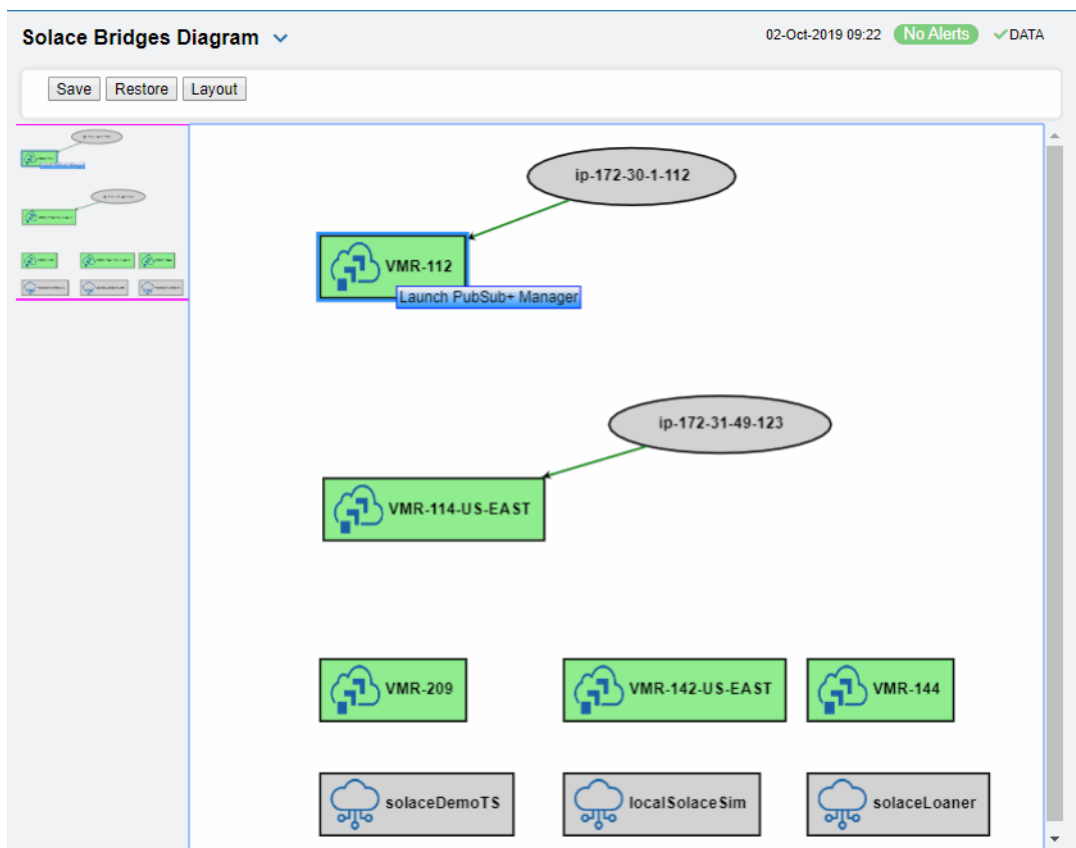
Restore: Returns objects to their previous positions.

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 [“CSPF Neighbors Diagram”](#).



Bridge Summary

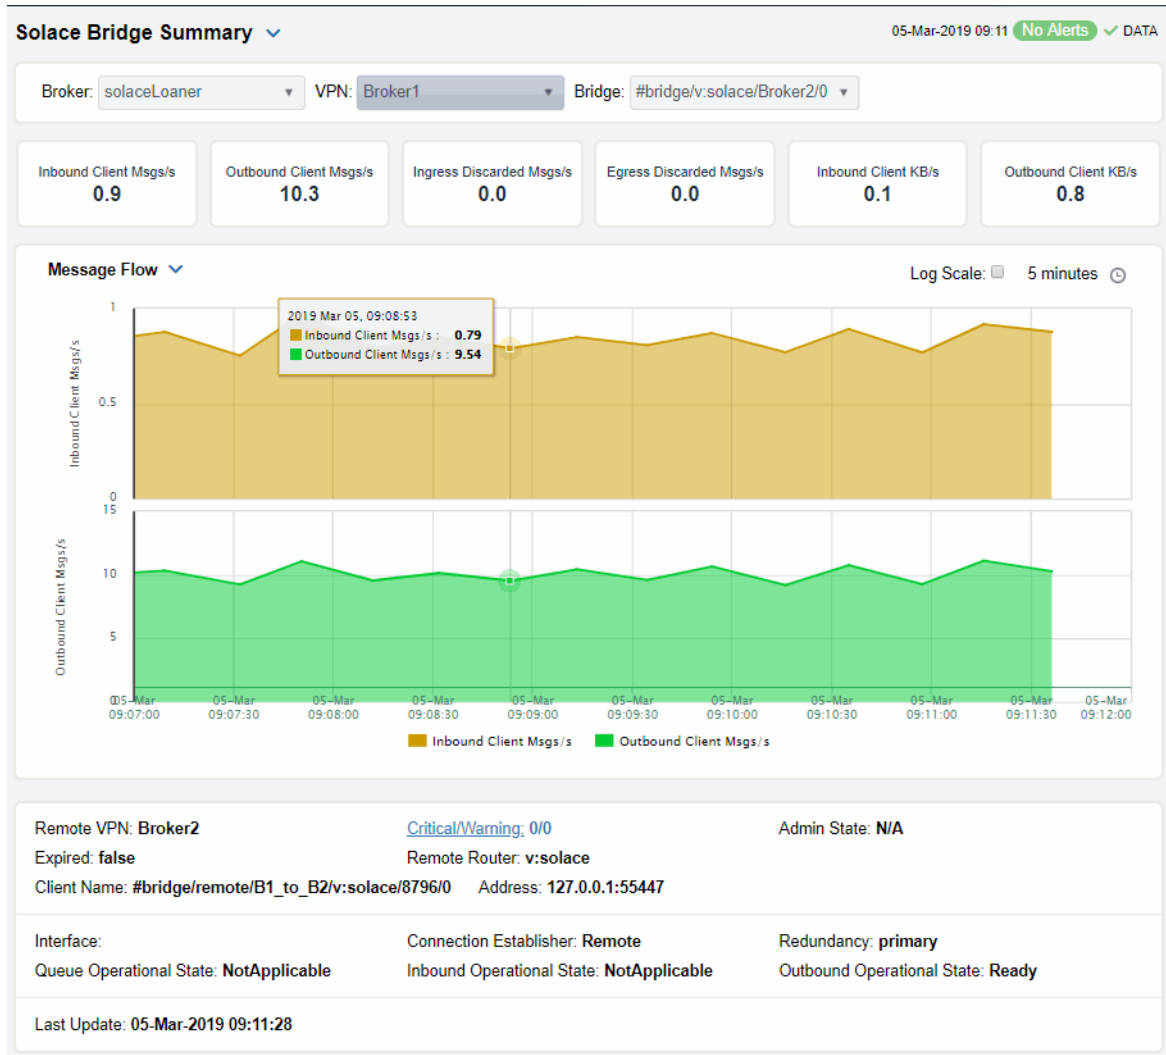
View current and historical performance and utilization metrics for a particular bridge on a VPN.

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.

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.

Conn Establisher

Refer to Solace documentation for more information.

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.)

Admin State

Indicates whether the bridge has been administratively enabled (via SolAdmin or the command line interface).

Client Name

The name of the client.

Redundancy	Indicates whether the bridge is the primary bridge, the backup bridge, the static bridge (default bridge used when no other bridge is available), or whether it is the only bridge available (none).
Outbound Op 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.)

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 , use the **Show** drop-down menu to include **All**, **Expired** or **Unexpired** and use the **Type** drop-down menu to see **All**, **Queues Only** or **Topic Endpoint Only**.

When queue and topic from the same broker and VPN share the same name, **-TE** is appended to the topic name.

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.

Broker	VPN	Endpoint Name	Alert Level	Q/T Type	Endpoint Type	Durable
solaceLoaner	BridgeBroker1	BidgeBroker1Queue	✓	Q	Primary	✓
solaceLoaner	Broker1	bridgeq	✓	Q	Primary	✓
solaceLoaner	Broker1	q1	✓	Q	Primary	✓
solaceLoaner	Broker10	Broker10Queue	✓	Q	Primary	✓
solaceLoaner	Broker10	Broker10Topic1	✓	T	Primary	✓
solaceLoaner	Broker2	broker2q1	✓	Q	Primary	✓
solaceLoaner	Broker2	msg_from_b1	✓	Q	Primary	✓
solaceLoaner	Broker3	Broker3Queue	✓	Q	Primary	✓
solaceLoaner	Broker4	Broker4Queue	✓	Q	Primary	✓
solaceLoaner	Broker5	Broker5Queue	✓	Q	Primary	✓
solaceLoaner	Broker6	Broker6Queue	✓	Q	Primary	✓
solaceLoaner	Broker7	Broker7Queue	✓	Q	Primary	✓
solaceLoaner	Broker8	Broker8Queue	✓	Q	Primary	✓
solaceLoaner	Broker9	Broker9Queue	✓	Q	Primary	✓

- Broker:** Select **All** brokers or a particular broker.
- VPN:** Select **All** VPNs or a particular VPN.
- Show Ingress Config Status Down Only** When checked, only shows endpoints configured on a VPN that currently have this status.
- Show:** Select to show **All**, **Expired** or **Unexpired** endpoints configured on the VPN(s).
- Type:** Select to show **All**, **Queues Only** or **Topic Endpoint Only** on the VPN(s).
- Filter Endpoint Name:** Enter a string to limit list in the table.
- Endpoints:** The number of rows currently in the table.
- VPN** The name of the VPN.
- Endpoint** The name of the endpoint.
- Alert Level** The current alert severity 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 endpoint.

Bind Count	The total number of binds connected to the endpoint.
Endpoint Type	The type of endpoint (either queue or topic).
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.
In Config Status	Refer to Solace documentation for more information.
Out Config Status	Refer to Solace documentation for more information.
Type	Refer to Solace documentation for more information.
Access Type	Refer to Solace documentation for more information.
Spooled Messages	The total number of spooled messages on the endpoint.
Spool Usage (MB)	The total spool usage consumed on the endpoint (in megabytes).
High Water Mark (MB)	The highest level of spool usage on the endpoint (in megabytes).
In Selector	Refer to Solace documentation for more information.
Out Selector	Refer to Solace documentation for more information.
Expired	When checked, performance data about the endpoint has not been received within the time specified.
Time Stamp	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 spooled 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.

**Spooled Messages**

The total number of spooled 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.

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** .

Endpoint Type

The type of endpoint.

Egress Config Status

Refer to Solace documentation for more information.

Egress Selector Present

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**.

Access Type

Refer to Solace documentation for more information.

Ingress Config Status

Refer to Solace documentation for more information.

Ingress Selector Present

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:

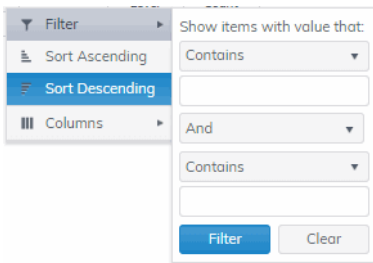
- **“Broker Capacity Table”**: View client, spool usage, incoming messages, outgoing messages, incoming bytes, and outgoing bytes data for all brokers.
- **“Broker Capacity - Summary”**: View client, spool usage, incoming messages, outgoing messages, incoming bytes, and outgoing bytes data for a specific broker.
- **“Broker Capacity Trends”**: View the broker capacity data for a specific broker in a trend graph format.

Broker 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:






Double-click a row to drill down and investigate in the **“Broker Capacity - Summary”** display.

Broker Capacity Table 11-Jun-2020 16:03 No Alerts DATA

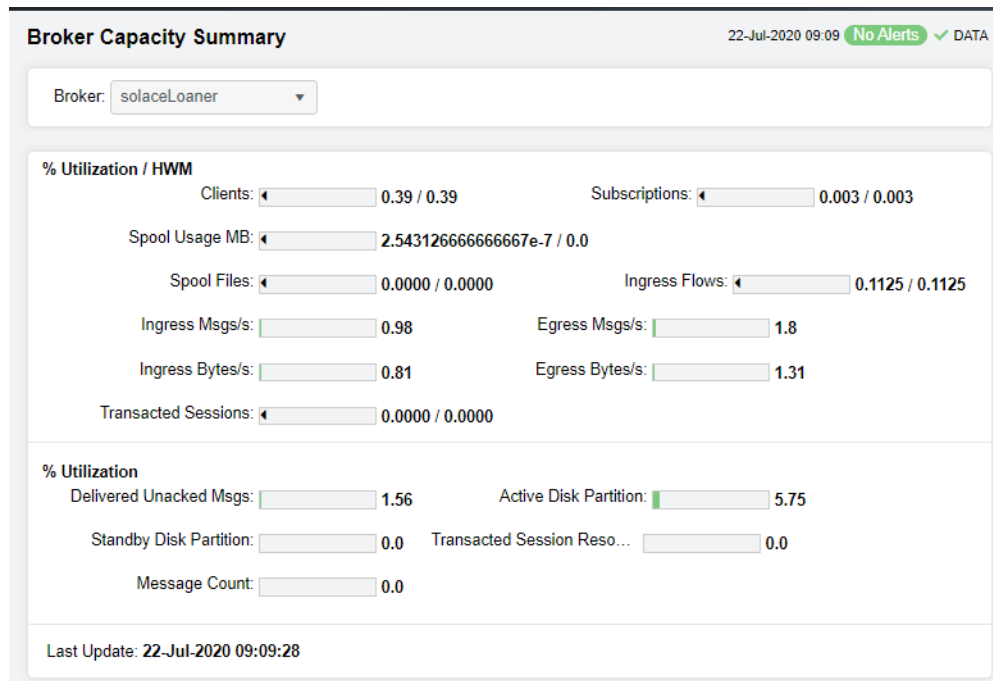
Show: All Brokers: 5 Fewer Columns

Broker	Alert Level	Alert Count	Current Client Connections	Connections HWM	Connections Max	Connections Used %	Connect Used HV
VMR_112-SP	✓		2	2	100	2.0	
VMR_114-US-EAST-SP	✓		3	3	100	3.0	
VMR_142-US-EAST-SP	✓		3	3	100	3.0	
VMR_144-SP	✓		4	4	100	4.0	
VMR_209	✓		3	3	100	3.0	

Broker	The name of the broker.
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.
Current Client Connections	The current number of clients connected.
Connections HWM	The greatest number of connections in the last 30 days.
Connections Max	The greatest number of connections since the broker last started.
Connections Used %	The current amount of connections used, in percent.
Connections Used HWM %	The greatest amount of connections used, in percent, in the last 30 days.
Cur Spool Usage MB	The current amount of used spool disk, in megabytes.
Cur Spool Usage HWM	The greatest amount of spool disk used in the last 30 days.
Spool Size Max MB	The amount of allocated spool disk.
Current Spool Usage %	The current amount of used spool disk, in percent.
Current Spool Usage % HWM	The greatest amount of used spool disk in the last 30 days, in percent.
Delivered Unacked Msgs Util %	Refer to Solace documentation for more information.
Ingress Flow Count	The number of ingress flows.
Ingress Flow HWM	The greatest number of ingress flows in the last 30 days.
Ingress Flows Allowed	The maximum number of ingress flows allowed.
Ingress Flow Count %	The amount of ingress flows in percent.
Ingress Flow Count HWM %	The greatest amount of ingress flows in the last 30 days, in percent.
Ingress Msgs/s	The number of ingress messages per second.
Ingress Msgs/s HWM	The greatest number of ingress messages per second in the last 30 days.
Cur Egress Msgs/s	The number of egress messages per second.
Egress Msgs/s HWM	The greatest number of egress messages per second in the last 30 days.
Egress Bytes/s HWM	The greatest amount of egress, in bytes per second, in the last 30 days, in percent.
Time Stamp	The date and time the row of data was last updated.

Broker Capacity - Summary

This display, a pivoted view of the “[Broker Capacity Table](#)”, allows you to view current and HWM (high water mark for the last 30 days) capacity utilization data for a single broker. Current values are represented by the green bar next to each field. 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. You can mouse over to see rate-per-second values for **Egress** and **Ingress Bytes/s** and **Egress** and **Ingress Msgs/s**.



% Utilization/HWM The following fields show two values: the current % (percent) utilization / peak capacity utilization (HWM) for the last 30 days. HWM values are represented by the black triangular pointer on a bar that reads from **0** to **100%**.

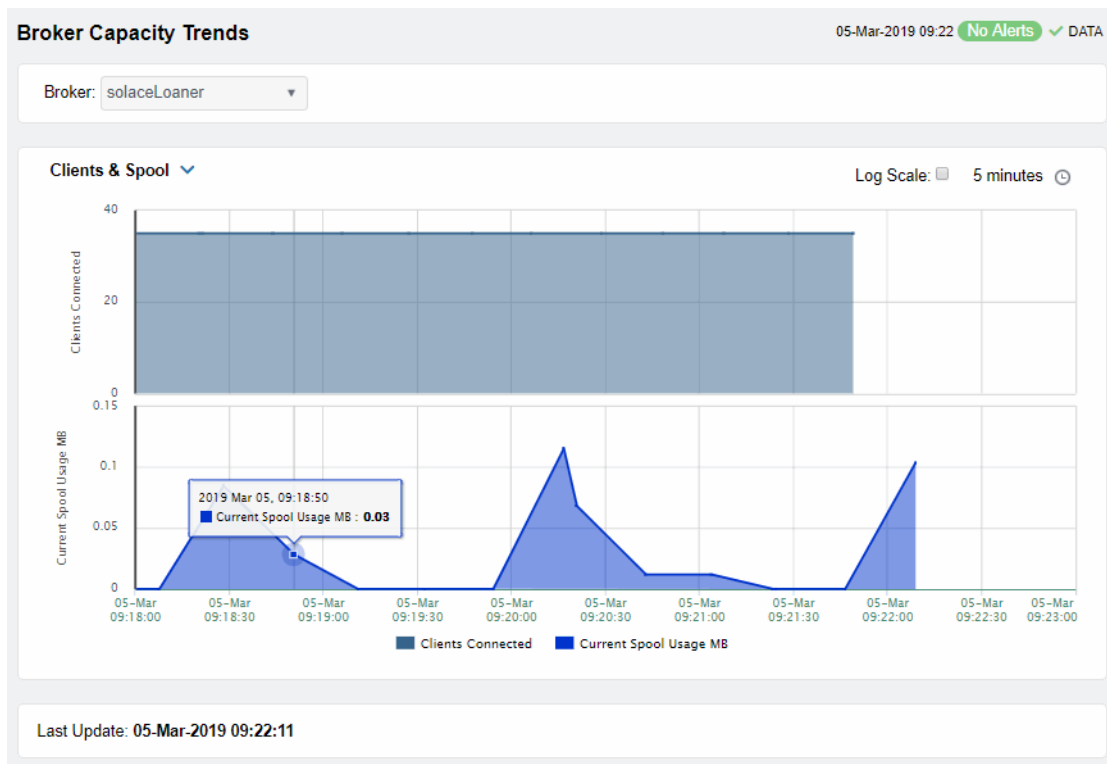
Clients	The clients connected to the broker.
Spool Files	The number of spool files on the broker.
Egress Msgs/s	The outgoing messages per second for the broker.
Transacted Sessions	The number of transacted sessions on the broker.
Subscriptions	The number of subscriptions on the broker.
Ingress Flows	The number of inflows on the broker.
Ingress Bytes/s	The incoming bytes per second for the broker.
Spool Usage MB	The amount of spool utilization, in megabytes per second, for the broker.
Ingress Msgs/s	The incoming messages per second for the broker.
Egress Bytes/s	The outgoing bytes per second for the broker.

% Utilization The following fields show current capacity % utilization.




Delivered Unacked Msgs	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.
Transacted Sessions Reso...	The current number of transacted sessions that were resolved on the broker.
Active Disk Partition	The percentage of available active disk partition that is used.
Message Count	The current number of messages on the broker.
Standby Disk Partition	The percentage of available standby disk partition that has been used.
Last Update	The date and time of the last data update.


Broker 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 displays allows you to supervise the last Syslog event messages from the Solace Message Brokers that have been configured for Syslog monitoring. See Solace product documentation for an in depth description of Syslog monitoring in Solace products and how to configure Message Brokers and the Syslog destination.

These displays require the Solace Event Module from the RTView Solace Monitor to be properly configured with a Syslog destination and running. Displays in this View are:

- **"Syslog Events Table"**: A tabular view of all available data for all Syslog events configured on a specific broker.
- **"Syslog Event Summary"**: Expanded/Summary information for a single Syslog event.

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. Click a row to investigate the Syslog event in the “[Syslog Event Summary](#)” display.

Search by clicking the right side of a column heading/**Filter** to open the Search, Sort and Choose Columns dialog:

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 Fi...	Cl
26-Mar-2019 12:00:00	VMR-47	SYSTEM	ip-172-30-1-144	19	SYSTEM_AUTH	notice	SYSLOG_MSG	Tue Mar 26 19:3	("sessionType":	
26-Mar-2019 12:00:00	VMR-118	SYSTEM	ip-172-30-1-112	19	SYSTEM_AUTH	notice	SYSLOG_MSG	Tue Mar 26 19:3	("sessionType":	
26-Mar-2019 12:00:00	VMR-47	SYSTEM	ip-172-30-1-144	19	SYSTEM_AUTH	notice	SYSLOG_MSG	Tue Mar 26 19:3	("sessionType":	
26-Mar-2019 12:00:00	VMR-118	SYSTEM	ip-172-30-1-112	19	SYSTEM_AUTH	notice	SYSLOG_MSG	Tue Mar 26 19:3	("sessionType":	
26-Mar-2019 12:00:00	VMR-47	SYSTEM	ip-172-30-1-144	19	SYSTEM_AUTH	notice	SYSLOG_MSG	Tue Mar 26 19:3	("sessionType":	
26-Mar-2019 12:00:00	VMR-118	SYSTEM	ip-172-30-1-112	19	SYSTEM_AUTH	notice	SYSLOG_MSG	Tue Mar 26 19:3	("sessionType":	
26-Mar-2019 12:00:00	VMR-47	SYSTEM	ip-172-30-1-144	19	SYSTEM_AUTH	notice	SYSLOG_MSG	Tue Mar 26 19:3	("sessionType":	
26-Mar-2019 12:00:00	VMR-118	SYSTEM	ip-172-30-1-112	19	SYSTEM_AUTH	notice	SYSLOG_MSG	Tue Mar 26 19:3	("sessionType":	
26-Mar-2019 12:00:00	VMR-47	SYSTEM	ip-172-30-1-144	19	SYSTEM_AUTH	notice	SYSLOG_MSG	Tue Mar 26 19:3	("sessionType":	
26-Mar-2019 12:00:00	VMR-118	SYSTEM	ip-172-30-1-112	19	SYSTEM_AUTH	notice	SYSLOG_MSG	Tue Mar 26 19:3	("sessionType":	
26-Mar-2019 12:00:00	VMR-47	SYSTEM	ip-172-30-1-144	19	SYSTEM_AUTH	notice	SYSLOG_MSG	Tue Mar 26 19:3	("sessionType":	
26-Mar-2019 12:00:00	VMR-118	SYSTEM	ip-172-30-1-112	19	SYSTEM_AUTH	notice	SYSLOG_MSG	Tue Mar 26 19:3	("sessionType":	
26-Mar-2019 12:00:00	VMR-47	SYSTEM	ip-172-30-1-144	19	SYSTEM_AUTH	notice	SYSLOG_MSG	Tue Mar 26 19:3	("sessionType":	
26-Mar-2019 12:00:00	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.

Scope:	<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"> • SYSTEM • VPN • CLIENT • ALL shows messages from all sources.
Severity:	<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"> • INFO • NOTICE • NOTICE or higher • WARN • WARN or higher • ERROR • ERROR or higher • CRITICAL • ALERT • EMERGENCY • ALL shows messages regardless of severity level from all sources.
Show:	<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"> • Expired Only • Unexpired Only • ALL shows expired and unexpired messages from all sources.
Events:	The number of events currently shown in the table.
Time Stamp	The date and time the row of data was last updated.

Syslog Event Summary

View status and configuration details for a single Syslog Event, such as the **Syslog Severity Code**, the **Intended Support Group**, and whether the event is **Clearable** and **Expired**.

Select a **Connection** and an **ID** from the drop-down menus.

Refer to Solace product documentation for additional information about values provided in this display.

Solace Syslog Event Summary 16-Jun-2020 14:48 No Alerts DATA

Connection: centos7-1
 ID: vm-centos7-1_SYSTEM_AD_MSG_CACHE_USAGE_HIGH_CLEAR_1592336901637

Connection: centos7-1 Router: vm-centos7-1 Solace Scope: SYSTEM Facility: 19 Host: vm-centos7-1

Solace Event ID: SYSTEM_AD_MSG_CACHE_USAGE_HIGH_CLEAR
 ID: vm-centos7-1_SYSTEM_AD_MSG_CACHE_USAGE_HIGH_CLEAR_1592336901637

Solace Event Source: event.log Syslog Severity: Informational Clearable: false
 Is Clear Event: true Syslog Severity Code: 6 Type: SYSLOG_MSG

Message:
 Tue Jun 16 10:54:16 PDT 2020 vm-centos7-1 event: SYSTEM: SYSTEM_AD_MSG_CACHE_USAGE_HIGH_CLEAR: -- Guaranteed Message Cache Usage Threshold 6.0 (60%) cleared: 5.2

Solace Description: Guaranteed Message Cache Usage Threshold 6.0 (60%) cleared: 5.2

Additional Fields: {"actual": 5,"threshold": 6,"thresholdPct": 60,"actualTenths": 2,"thresholdTenths": 0}

Clearable: false Is Clear Event: true Intended Support Group: APPLICATION
 Solace Vpn: - Solace Client: - Solace Company: -
 Solace Application: - Solace Environment: - Expired: false

Last Update: 16-Jun-2020 14:48:21

Connection:	Select a connection.
ID:	Select an event ID.
Connection	The name of the selected connection.
Router	The name of the router for the selected event.
Solace Scope	The type of Syslog event: <ul style="list-style-type: none"> • 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.
Facility	The generating facility for the event. Event broker Syslog messages belong to a facility, which is a group of messages that are either generated by the same software process, or concern a similar event broker subsystem condition or activity (such as debugging attempts).
Host	The name of the host associated with the event.
Solace Event ID	The type of Solace event.
ID	The unique identifier for the event.

Solace Event Source	The name of the log file in which the event is recorded.
Is Clear Event	true/false
Syslog Severity:	The Syslog severity level for the event: 0 - Emergency: system is unusable. 1 - Alert: action must be taken immediately. 2 - Critical: critical conditions. 3 - Error: error conditions. 4 - Warning: warning conditions. 5 - Notice: normal but significant condition. 6 - Informational: informational messages. 7 - Debug: debug-level messages.
Syslog Severity Code:	0 - Emergency: system is unusable. 1 - Alert: action must be taken immediately. 2 - Critical: critical conditions. 3 - Error: error conditions. 4 - Warning: warning conditions. 5 - Notice: normal but significant condition. 6 - Informational: informational messages. 7 - Debug: debug-level messages.
Clearable:	true/false
Type:	The type of the selected event.
Message:	The Syslog message content.
Solace Description	The Message Text for this event.
Additional Fields	Additional information associated with the event (formatted as Name:Value pairs in a JSON object).
Clearable	true/false
Solace Vpn	The Solace VPN identifier associated with this event.
Solace Application	The Solace Application identifier associated with this event.
Is Clear Event	true/false
Solace Client	The Solace Client identifier associated with this event.
Solace Environment	The Solace environment identifier associated with this event .
Intended Support Group	The intended support group for this event.
Solace Company	The Solace Company identifier associated with this event.
Expired	true/false
Last Update	The date and time of the last data update.

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.

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 \(2.x\)"](#)
- ["TIBCO Adapters"](#)
- ["TIBCO BusinessEvents"](#)
- ["TIBCO BusinessWorks"](#)
- ["TIBCO BusinessWorks 5 Monitor"](#)
- ["TIBCO Enterprise Message Service"](#)
- ["TIBCO FTL"](#)
- 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.

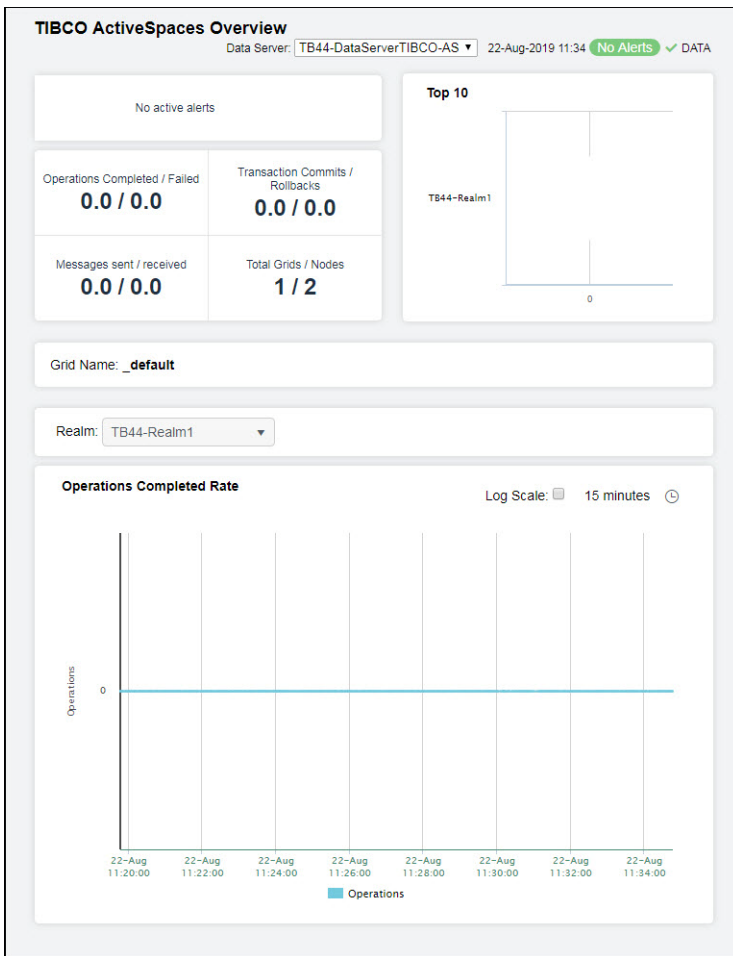
TIBCO ActiveSpaces

The **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 grids and nodes.
- A visual list of the top 10 realms containing the best operations completed rate 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 the **Operations Completed Rate** trend graph representing Operations Completed Rate 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**:

- **"Grids Views"**: The displays in this View provide detailed data for all grids in a heatmap and tabular format, or for a particular grid in tabular and trend graph format.
- **"Nodes Views"**: The displays in this View provide detailed data for all nodes in a heatmap or tabular format.
- **"Proxies Views"**: The displays in this View provide detailed data for all proxies in a heatmap and tabular format.
- **"Keepers Views"**: The displays in this View provide detailed data for all keepers in a heatmap or tabular format.

Grids Views

These displays provide detailed data for all grids in a heatmap and tabular format. Clicking **Grids** in the left/navigation menu opens the ["TIBCO ActiveSpaces Grids Table"](#), which provides a tabular view of your grids and their associated metrics. Displays in this View are:

- **All Grids Heatmap:** Opens the ["TIBCO ActiveSpaces Grids Heatmap"](#) display, which provides a heatmap view of all grids.
- **Single Grid Summary:** Opens the ["TIBCO ActiveSpaces Grid Summary"](#) display, which provides a view of the current and historical metrics for a single grid.
- **Realm Servers:** Opens the ["TIBCO ActiveSpaces Realm Server"](#) display, which provides a view of the server CPU percent usage and the server memory used by the realm server managing the grid (in KBs) in a trend graph format.




TIBCO ActiveSpaces Grids Table

The table in this display provides a view of all of your grids 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 grid in the ["TIBCO ActiveSpaces Grid Summary"](#) display.

Grid Name	Realm	Alert Level	Alert Count	Expired	Server Version	Ops Completed	Cc
_default	TB44-Realm1	✓			6.1.0 V5	0	

All Grids Table

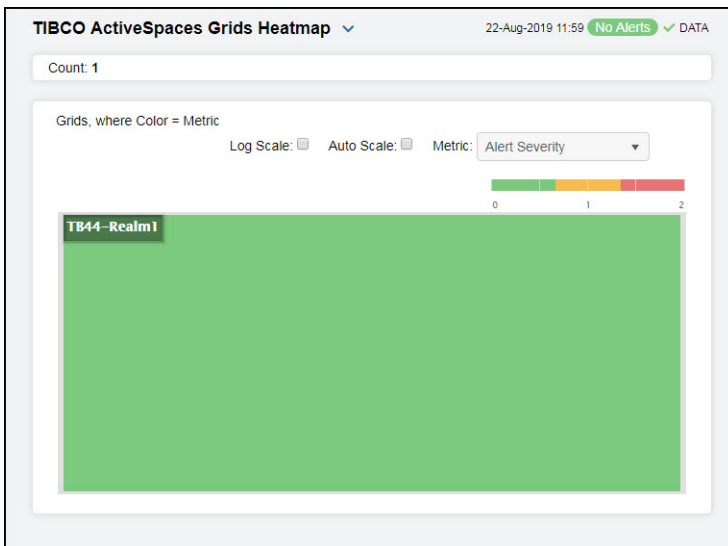
Grid Name	The name of the grid.
Realm	The name of the realm.

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.
Ops Completed	The total number of operations completed in the grid.*
Ops Completed Rate	The rate of operations completed in the grid.*
Ops Failed	The total number of operations failed in the grid.*
Ops Failed Rate	The rate of operations failed in the grid.*
Node Live Data	The size of the node's live data.*
Txn Commits	The number of transaction commits in the grid.*
Txn Commits Rate	The rate of transaction commits in the grid.*
Txn Rollbacks	The number of transaction rollbacks in the grid.*
Txn Rollbacks Rate	The rate of transaction rollbacks in the grid.*

TIBCO ActiveSpaces Grids Heatmap

Clicking **All Grids Heatmap** in the left/navigation menu opens the **TIBCO ActiveSpaces Grids Heatmap**, which provides an easy-to-view interface that allows you to quickly identify the current status of each of your grids for each available metric. You can view the grids in the heatmap based on the following metrics: current alert severity, alert count, 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 grid. The rectangle color indicates the most critical alert state. Click on a node to drill-down to the "**TIBCO ActiveSpaces Grid Summary**" display and view metrics for a particular grid. Toggle between the commonly accessed displays by clicking the drop down list on the display title. Mouse-over rectangles to view more details about grid performance and status.




Fields and Data:




Log 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** 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.

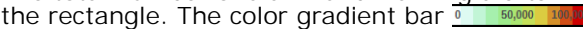
Auto Scale When checked, the values of the selected metric are auto-scaled to its highest defined value. When unchecked, the values of the selected metric display based on the threshold defined for the alert associated with the selected metric. Selecting Auto helps to visualize the range of the values 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

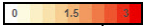
Select the metric driving the heatmap display. The default is **Alert Severity**. Each **Metric** has a color gradient bar that maps values to colors. The heatmap is organized by grids, where each rectangle represents a grid. Mouse-over any rectangle to display the current values of the metrics for the grid. Click on a rectangle to drill-down to the associated ["TIBCO ActiveSpaces Grid Summary"](#) display for a detailed view of metrics for that particular grid.

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.

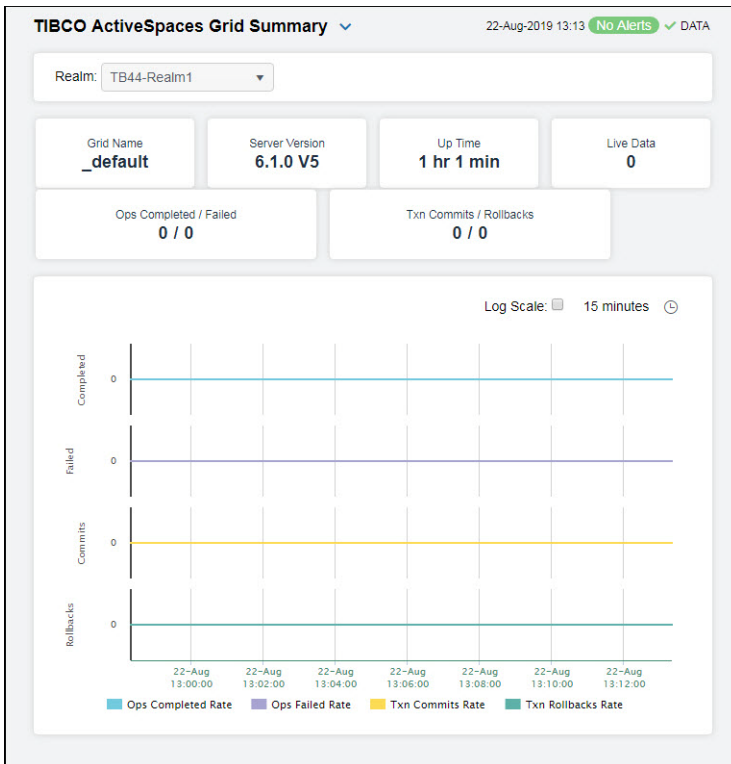
Alert Count The total number of alarm and warning alerts in a given item (index) associated with the rectangle. The color gradient bar  shows the range of the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

Ops Completed The number of operations completed in the grid. The color 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 grid. 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 Grid Summary

Clicking **Single Grid Summary** in the left/navigation menu opens the **TIBCO ActiveSpaces Grid Summary** display, which provides a view of the current and historical metrics for a single grid. 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.



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:

Grid Name The name of the grid.

Server Version The version of the realm server.

Up Time The amount of time since the realm server was started.

Live Data The size of the live data.

**Ops Completed/
Failed** The number of operations completed and the number of operations failed in the grid.*

**Txn Commits/
Rollbacks** The total number of transactions committed and the number of transactions rolled back in the grid.*

Trends Traces the following:

Ops Completed Rate -- traces the number of operations completed per second.

Ops Failed Rate-- traces the number of operations failed per second.

Txn Commits Rate -- traces the number of transactions committed per second.


Txn Rollbacks Rate -- 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.

Time Settings

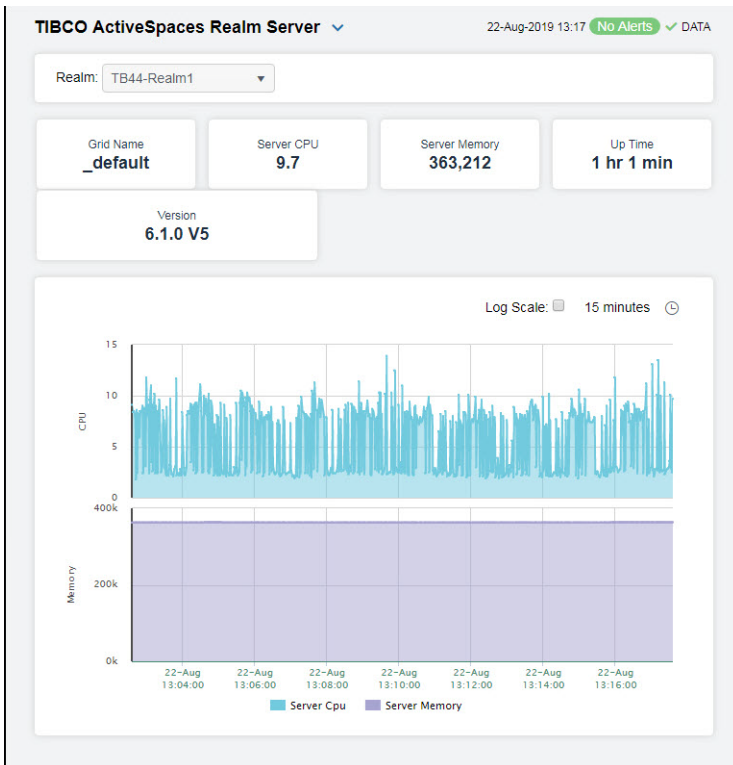
Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

TIBCO ActiveSpaces Realm Server

Clicking **Realm Server** in the left/navigation menu opens the **TIBCO ActiveSpaces Realm Server** display, which provides a view of the current and historical metrics for the realm server of a single grid. Hover over the boxes at the top of the display to view additional information. In the trend graph region, you can view the server CPU percent usage and the server memory used (in KBs).

**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:

Grid Name The name of the grid.

Server CPU The server's CPU usage percentage.

Server Memory The used memory on the server, in kilobytes.

Up Time The amount of time since the server was started.

Version The version of the server.


Trends

Traces the following:

Server CPU -- traces the server's CPU usage percentage.

Server Memory -- traces the used memory on the 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.
- Time Settings** Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

Nodes Views

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"](#), which provides a tabular view of all nodes (contained within a particular grid) and their associated metrics. Displays in this View are:

- **All Nodes Heatmap**: Opens the ["TIBCO ActiveSpaces Nodes Heatmap"](#) display, which is a heatmap view of all nodes contained within a particular grid.
- **Single Node Summary**: Opens the ["TIBCO ActiveSpaces Node Summary"](#) display, which allows you to view metrics and trend data for a particular node.

TIBCO ActiveSpaces Nodes Table

The table in this display provides a view of all nodes and their associated metric data in a specific grid. 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”](#) display.

Grid Name	Copyset	Node	Realm	Alert Level	Alert Count	Expired	CPI Used
__default	set1	n1	TB44-Realm1	Green			
__default	set2	n2	TB44-Realm1	Green			

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 nodes found in the search and listed in the table.

All Nodes Table

Grid Name The name of the grid associated with the node.

Copyset The name of the Copyset hosted by the node.*

Node The name of the node.

Realm The name of the realm in which the node resides.

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 ID of the node.*

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.

CPU Used (%)	The percentage of CPU memory used on the node.*
Memory Used (kb)	The amount of memory used by the node, in kilobytes.
Client Status	The current status of the client on which the node resides*
Bytes Received	The number of bytes received.
Bytes Received/s	The rate of bytes received.
Bytes Sent	The number of bytes sent.
Bytes Sent/s	The rate of bytes sent.
Messages Received	The number of message received.
Messages Rcvd Rate	The rate of messages received on the node.
Messages Sent	The number of messages sent.
Messages Sent Rate	The rate of messages sent on the node.
Live Data Size	The size of the live data.*
Put Rate	The rate of "put" operations (per second) performed on the node.*
Get Rate	The rate of "get" operations (per second) performed on the node.*
Remove Rate	The rate of "remove" operations (per second) performed on the node.*
Client ID	The ID of the client.*
Host	The name of the host.*
PID	The process ID of the node process.*
Node Ready	When checked, the node is operational.
Node Started	When checked, the node has been started and is up and running.*
Get Count	The total number of "get" operation performed on the node.*
Put Count	The total number of "put" operations performed on the node.*
Remove Count	The total number of "remove" operations performed on the node.*
Transaction Begin Rate	The rate of transactions started on the node.*
Transaction Commit Rate	The rate of transactions committed on the node.*
Transaction Rollback Rate	The rate of transactions that have been rolled back on the node.*

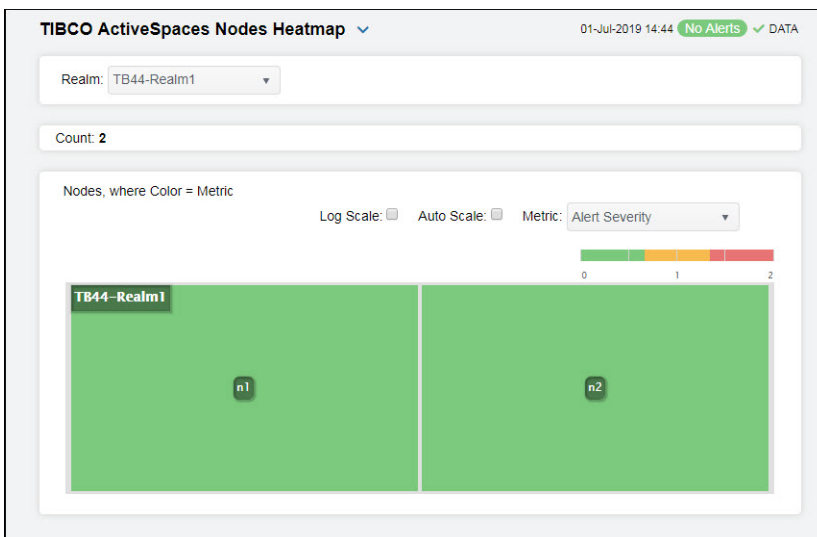
Iterator Create Rate	The rate of iterator operations being created on the node.*
Iterator Get Rate	The rate of "get" iterator operations on the node.*
Iterator Close Rate	The rate of iterator operations (per second) on the node.*
Query Create Rate	The rate of created queries on the node.*
Query Get Rate	The rate of "get" queries (per second) performed on the node.*
Query Close Rate	The rate of closed queries on the node.*
Operations Completed Rate	The rate at which operations are completed on the node.*
Operations Failed Rate	The rate at which operations are failing on the node.*
Operations Completed	The total number of operations completed on the node.*
Operations Failed	The number of operations failed on the node.*
Transaction Begin Count	The number of transactions started on the node.*
Transaction Commit Count	The number of transactions committed on the node.*
Operations Suspended	The number of operations suspended on the node.*
Transaction Rollbacks	The number of transactions that have been rolled back on the node.*
Iterator Creates	The number of iterator operations created on the node.*
Query Creates	The number of created queries on the node.*
Reindex Completed	The number of "reindex" operations completed on the node.*
Scans Completed	The number of completed scans on the node.*
Iterator Gets	The number of "get" iterator operations on the node.*
Query Gets	The total number of "get" queries performed on the node.*
Reindex Rows	The number of "reindex" rows on the node.*
Expired Rows	The number of expired rows on the node.*
Iterator Closes	The number of closed iterator operations on the node.*
Query Closes	The number of queries that were closed on the node.*

Operations Allowed	The number of operations allowed on the node.*
Operations Rejected	The number of operations rejected on the node.*
Operations Pending	The number of operations pending on the node.*
Host	The name of the host.*
Client Status	The current status of the client.*
Client ID	The ID of the client.*
PID	The process ID of the client.
Time Stamp	The date and time the row data was last updated.

TIBCO ActiveSpaces Nodes Heatmap

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”](#) display, which allows you to see additional details for the selected node.



Filter By:

Realm Select the realm for which you want to see data.

Fields and Data:


Count The number of nodes listed in the heatmap.




Log 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** 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.

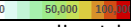
Auto Scale When checked, the values of the selected metric are auto-scaled to its highest defined value. When unchecked, the values of the selected metric display based on the threshold defined for the alert associated with the selected metric. Selecting Auto helps to visualize the range of the values 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).

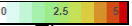
Metric

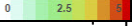
Select the metric driving the heatmap display. The default is **Alert Severity**. Each **Metric** has a color gradient bar that maps values to colors. The heatmap is organized by nodes, where each rectangle represents a node. Mouse-over any rectangle to display the current values of the metrics for the node. Click on a rectangle to drill-down to the associated ["TIBCO ActiveSpaces Node Summary"](#) display for a detailed view of metrics for that particular node.

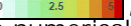
Alert Severity The current alert severity. Values range from 0 - 2, as indicated in the color gradient  bar, where 2 is the highest Alert Severity:


-  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  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 **TdgNodeCpuUsageHigh**. The middle value in the gradient bar indicates the middle value of the range.

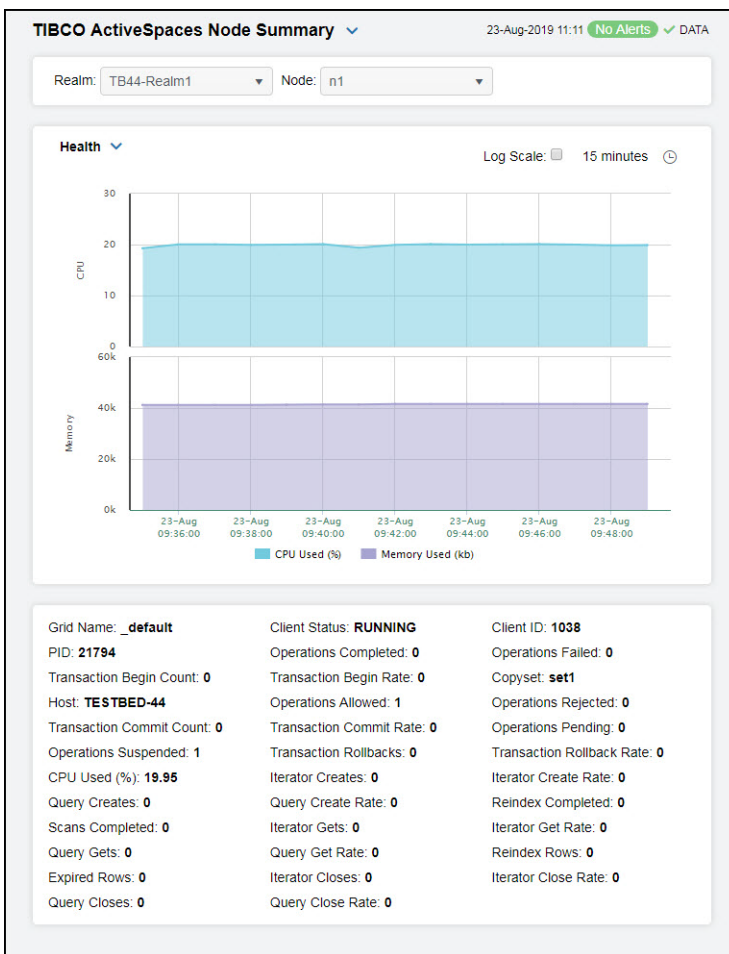
Memory The memory usage for the node. The color gradient bar  , populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **TdgNodeMemoryUseHigh**. The middle value in the gradient bar indicates the middle value of the range.

Failed Op Rate 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 **TdgNodeOpsFailedRateHigh**. The middle value in the gradient bar indicates the middle value of the range.

Completed Op Rate 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 **TdgNodeOpsCompletedRateLow**. The middle value in the gradient bar indicates the middle value of the range.

TIBCO ActiveSpaces Node Summary

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.




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.
- Performance Metric Trends**
- Health**
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**
Traces the following:
Live Data Size-- traces the Live Data Size.
- Operations**
Traces the following:
Operations Completed Rate -- traces the rate of completed operations.
Operations Failed Rate -- 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 Settings** Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

Fields and Data:

Grid Name	The name of the grid.*
PID	The process ID of the node.
Transaction Begin Count	The number of transactions started on the node.
Host	The name of the host.
Transaction Commit Count	The number of transactions committed on the node.
Operations Suspended	The number of suspended operations on the node.
CPU Used (%)	The percentage of CPU used
Query Creates	The number of created queries on the node.
Scans Completed	The number of scans completed.
Query Gets	The number of "get" operations on the node.
Expired Rows	The number of expired rows on the node.
Query Closes	The number of closed queries on the node.
Client Status	The current status of the node.
Operations Completed	The number of completed operations on the node.
Transaction Begin Rate	The rate of transactions started on the node.
Operations Allowed	The number of allowed operations on the node.
Transaction Commit Rate	The rate of transactions committed on the node.
Transaction Rollbacks	The number of transactions that have been rolled back on the node.
Iterator Creates	The number of iterator operations on the node.
Query Create Rate	The rate of created queries on the node.
Iterator Gets	The number of "get" iterator operations on the node.
Query Get Rate	The rate of "get" operations on the node.
Iterator Closes	The number of closed iterator operations on the node.
Query Close Rate	The rate of closed queries on the node.

Client ID	The ID of the node.*
Operations Failed	The number of failed operations on the node.
Copysset	The name of the copysset.*
Operations Rejected	The number of rejected operations on the node.
Operations Pending	The number of pending operations on the node.
Transaction Rollback Rate	The rate of transactions that have been rolled back on the node.
Iterator Create Rate	The rate of iterator operations on the node.
Reindex Completed	The number of "reindex" scans on the node.
Iterator Get Rate	The rate of "get" iterator operations on the node.
Reindex Rows	The number of "reindex" rows on the node.
Iterator Close Rate	The rate of closed iterator operations on the node.

Proxies Views

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](#)" 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](#)" display, which provides a heatmap view of all proxies contained within a particular realm.
- **Single Proxy Summary:** Opens the "[TIBCO ActiveSpaces Proxy Summary](#)" display, which allows you to view metrics and trend data for a particular proxy.

TIBCO ActiveSpaces 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 double-clicking a row to view details for the selected proxy in the "[TIBCO ActiveSpaces Proxy Summary](#)" display

TIBCO ActiveSpaces Proxies Table 22-Aug-2019 13:56 No Alerts DATA

Realm: TB44-Realm1

Count: 1

All Proxies Table

Grid Name	Proxy	Realm	Alert Level	Alert Count	Expired	CPU Used (%)	Memory Used (kb)
_default	p1	TB44-Realm1	Green			4.098	

Filter By:

Realm Select the realm for which you want to view data.

Count The total number of proxies found for the realm selected in the **Realm** dropdown, which are displayed in the **All Proxies Table**.

All Proxies Table

Grid Name The name of the grid.

Proxy The name of the proxy.

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 proxy.

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.

CPU Used (%) The percentage of CPU used on the proxy.*

Memory Used (kb) The amount of memory used, in kilobytes.*

Bytes Received The number of bytes received.

Bytes Received/s The rate of bytes received.

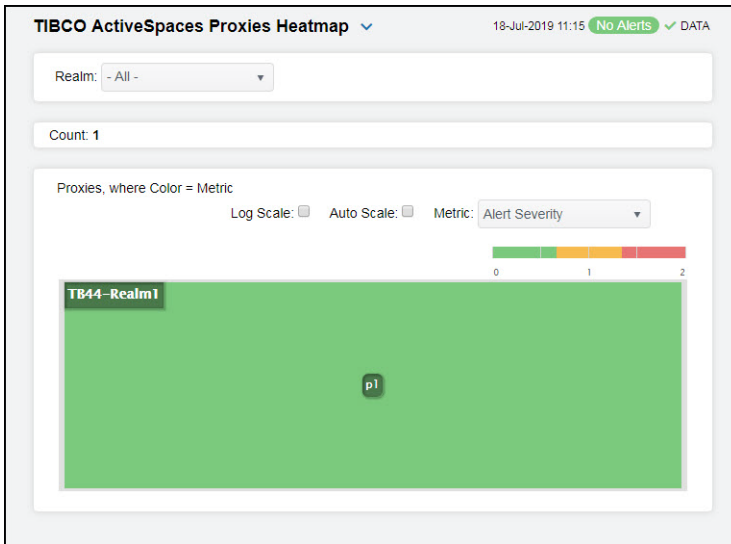
Bytes Sent	The number of bytes sent.
Bytes Sent/s	The rate of bytes sent.
Messages Received	The number of messages received.
Messages Received Rate	The rate of messages received.
Messages Sent	The number of messages sent.
Messages Sent Rate	The rate of messages sent.
Put Rate	The rate of "put" operations (per second) performed on the proxy.*
Get Rate	The rate of "get" operations (per second) performed on the proxy.*
Remove Rate	The rate of "remove" operations (per second) performed on the proxy.*
Transaction Begin Rate	The rate of transactions being started on the proxy.*
Transaction Commit Rate	The rate of transactions being committed on the proxy.*
Transaction Rollback Rate	The rate of transactions rolled back on the proxy.*
Iterator Create Rate	The rate of iterator operations being created on the proxy.*
Iterator Get Rate	The rate of "get" iterator operations on the proxy.*
Iterator Close Rate	The rate of iterator operations being closed on the proxy.*
Query Create Rate	The rate of created queries on the proxy.*
Query Get Rate	The rate of "get" queries on the proxy.*
Query Close Rate	The rate of closed queries on the proxy.*
Queries	The number of queries on the proxy.*
Statements	The number of statements on the proxy.*
Iterators	The number of Iterators on the proxy.*
Listeners	The number of listeners on the proxy.*
Client Status	The status of the client.*
Client ID	The ID of the client.*
PID	The process ID of the host.*
Gets	The total number of "get" operations performed on the proxy.*
Transaction Begins	The number of transactions started on the proxy.*

Client Connected	The number of clients connected.*
Host	The name of the host.*
Puts	The total number of "put" operations performed on the proxy.*
Transaction Commits	The number of commit transactions on the proxy.*
Removes	The total number of "remove" operations performed on the proxy.*
Transaction Rollbacks	The number of transactions rolled back on the proxy.*
Iterator Creates	The number of iterator operations on the proxy.*
Query Creates	The number of created queries on the proxy.*
Iterator Gets	The number of "get" iterator operations on the proxy.*
Query Gets	The number of "get" queries on the proxy..*
Iterator Closes	The number of iterator operations being closed on the proxy.*
Query Closes	The number of closed queries on the proxy.*
Time Stamp	The date and time the row data was last updated.

TIBCO ActiveSpaces Proxies Heatmap

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](#)" 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.




Log 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.

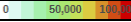
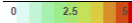

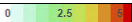
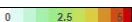
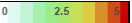
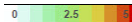
Auto Scale When checked, the values of the selected metric are auto-scaled to its highest defined value. When unchecked, the values of the selected metric display based on the threshold defined for the alert associated with the selected metric. Selecting Auto helps to visualize the range of the values 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 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"](#) display for a detailed view of metrics for that particular proxy.

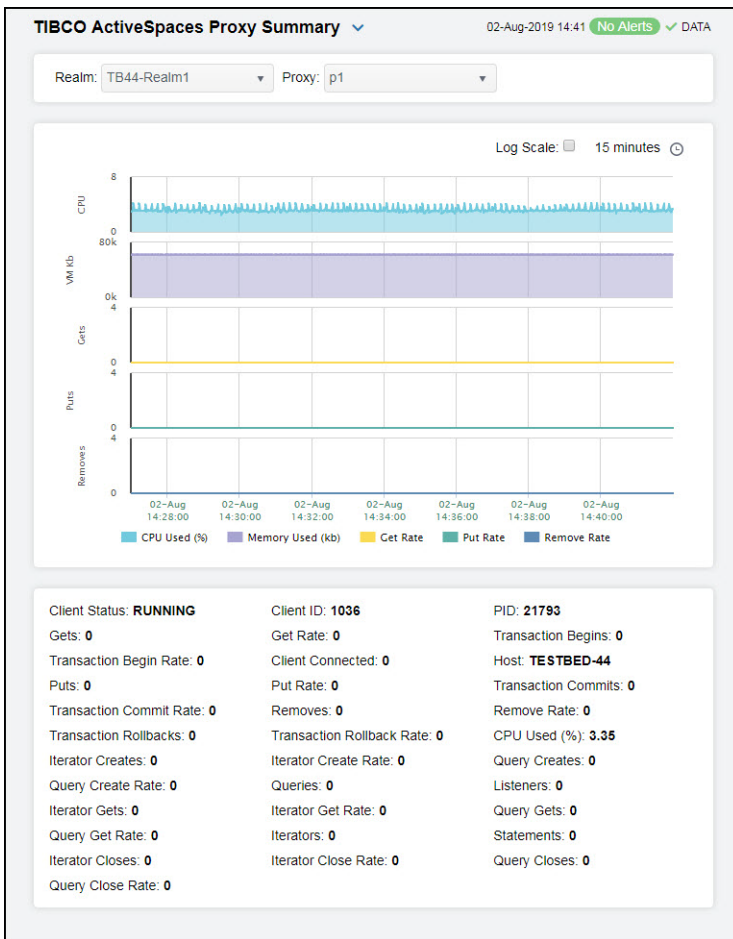
Alert Severity The current alert severity. Values range from **0 - 2**, as indicated in the color gradient  bar, where **2** is the highest Alert Severity:

-  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


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.

**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.
- 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 Settings** Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

Fields and Data:

Grid Name	The name of the grid.*
PID	The process ID of the proxy.
Transaction Begins	The number of transactions started on the proxy.
Host	The name of the host.
Transaction Commits	The number of transactions committed on the proxy.
Remove Rate	The rate of "remove" operations on the proxy.
CPU Used (%)	The percentage of CPU used.
Query Creates	The number of queries created on the proxy.
Listeners	The total number of listeners on the proxy.
Query Gets	The number of "get" queries on the proxy.

Statements	The total number of statements on the proxy.
Query Closes	The number of queries closed on the proxy.
Client Status	The current status of the proxy.
Gets	The number of "get" operations on the proxy.
Transaction Begin Rate	The rate of transactions being started on the proxy.
Puts	The number of "put" operations on the proxy.
Transaction Commit Rate	The rate of transactions being committed on the proxy.
Transaction Rollbacks	The number of transactions rolled back on the proxy.
Iterator Creates	The number of iterator operations created on the proxy.
Query Create Rate	The rate of queries being created on the proxy.
Iterator Gets	The number of "get" iterator operations on the proxy.
Query Get Rate	The rate of "get" queries being created on the proxy.
Iterator Closes	The number of closed iterator operations on the proxy.
Query Close Rate	The rate of queries being closed on the proxy.
Client ID	The ID of the proxy.*
Get Rate	The rate of "get" operations on the proxy.
Client Connected	The number of clients connected.
Put Rate	The rate of "put" operations on the proxy.
Removes	The number of "remove" operations on the proxy.
Transaction Rollback Rate	The rate of transactions being rolled back on the proxy.
Iterator Create Rate	The rate of iterator operations being created on the proxy.
Queries	The number of queries created on the proxy.
Iterator Get Rate	The rate of "get" iterator operations being created on the proxy.
Iterators	The number of iterator operations created on the proxy.
Iterator Close Rate	The rate of iterator operations being closed on the proxy.

Keepers Views

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](#)” 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](#)” display, which is a heatmap view of all keepers contained within a particular realm.
- **Single Keeper Summary:** Opens the “[TIBCO ActiveSpaces Keeper Summary](#)” display, which allows you to view metrics and trend data for a particular keeper.

TIBCO ActiveSpaces StateKeepers 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 double-clicking a row to view details for the selected keeper in the “[TIBCO ActiveSpaces Keeper Summary](#)” display

Grid Name	Keeper	Realm	Alert Level	Alert Count	Expired	CPU Used (%)	Mem Used (%)
_default	k1	TB44-Realm1	✓			2.499	

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

Grid Name The name of the grid.

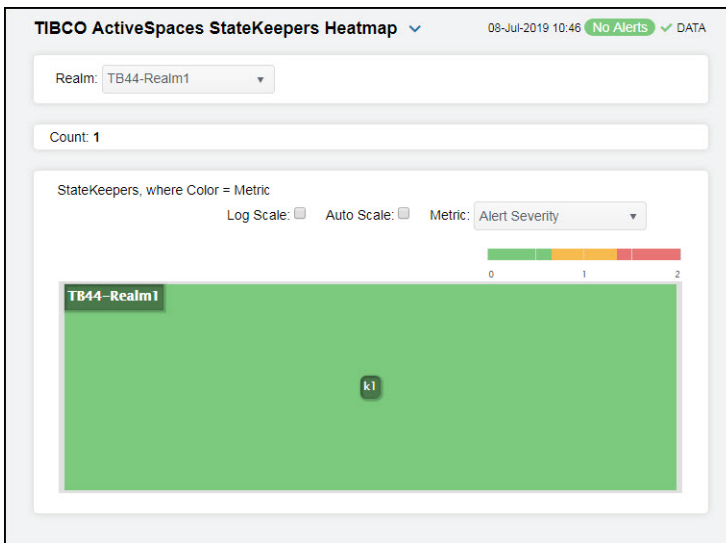
Keeper The name of the keeper.

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.
CPU Used (%)	The percentage of CPU memory used by the keeper.
Memory Used (kb)	The memory used by the keeper, in kilobytes.
Bytes Received	The number of bytes received.
Bytes Received/s	The rate of bytes received.
Bytes Sent	The number of bytes sent.
Bytes Sent/s	The rate of bytes sent.
Messages Received	The number of messages received.
Messages Rcvd Rate	The rate of messages received.
Message Sent Rate	The rate of messages sent.
Messages Sent	The number of messages sent.
Client Status	The current status of the client on which the keeper resides.*
Client ID	The ID of the client.*
PID	The process ID of the StateKeeper process.*
Host	The name of the host.*
Ready	When checked, the keeper is operational.*
Started	When checked, the keeper has been started and is up and running.*
Copyset Epoch Updated	Any value greater than 0 denotes that a disaster recovery failover to another data grid has occurred.
Time Stamp	The date and time the row data was last updated.

TIBCO ActiveSpaces StateKeepers Heatmap

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”](#) display, which allows you to see additional details for the selected keeper.



Filter By:

Realm Select the realm for which you want to see data.

Fields and Data:





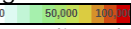

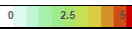
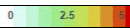
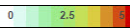
Count The total number of keepers found for the selected realm.

Log 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.

Auto Scale 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).

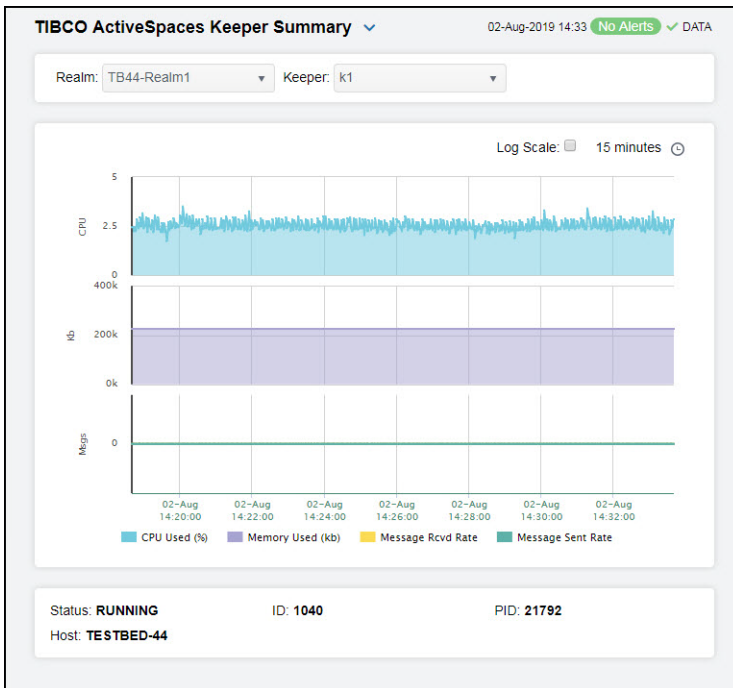
Metric

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 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"](#) display for a detailed view of metrics for that particular keeper.

Alert Severity	<p>The current alert severity. Values range from 0 - 2, as indicated in the color gradient  bar, where 2 is the highest Alert Severity:</p> <ul style="list-style-type: none">  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	<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>
CPU Usage	<p>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.</p>
Memory	<p>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.</p>
Msgs Rcvd/sec	<p>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.</p>
Msgs Sent/sec	<p>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.</p>

TIBCO ActiveSpaces Keeper Summary

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 percentage, process memory usage (in KB), rate of received messages, and the rate of sent messages for the keeper.



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.

Trends Traces the following:

CPU Used (%)-- traces the CPU usage percentage.


Memory Used (kb) -- traces the memory usage, in kilobytes.

Message Rcvd Rate-- traces the rate of messages received, per second.

Message Sent Rate-- traces the rate of 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 Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

Fields and Data:

Grid Name	The name of the grid.*
PID	The process ID of the StateKeeper.*
Status	The current status of the keeper.*
Host	The name of the host.
ID	The ID of the keeper.*

TIBCO ActiveSpaces (2.x)

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**:

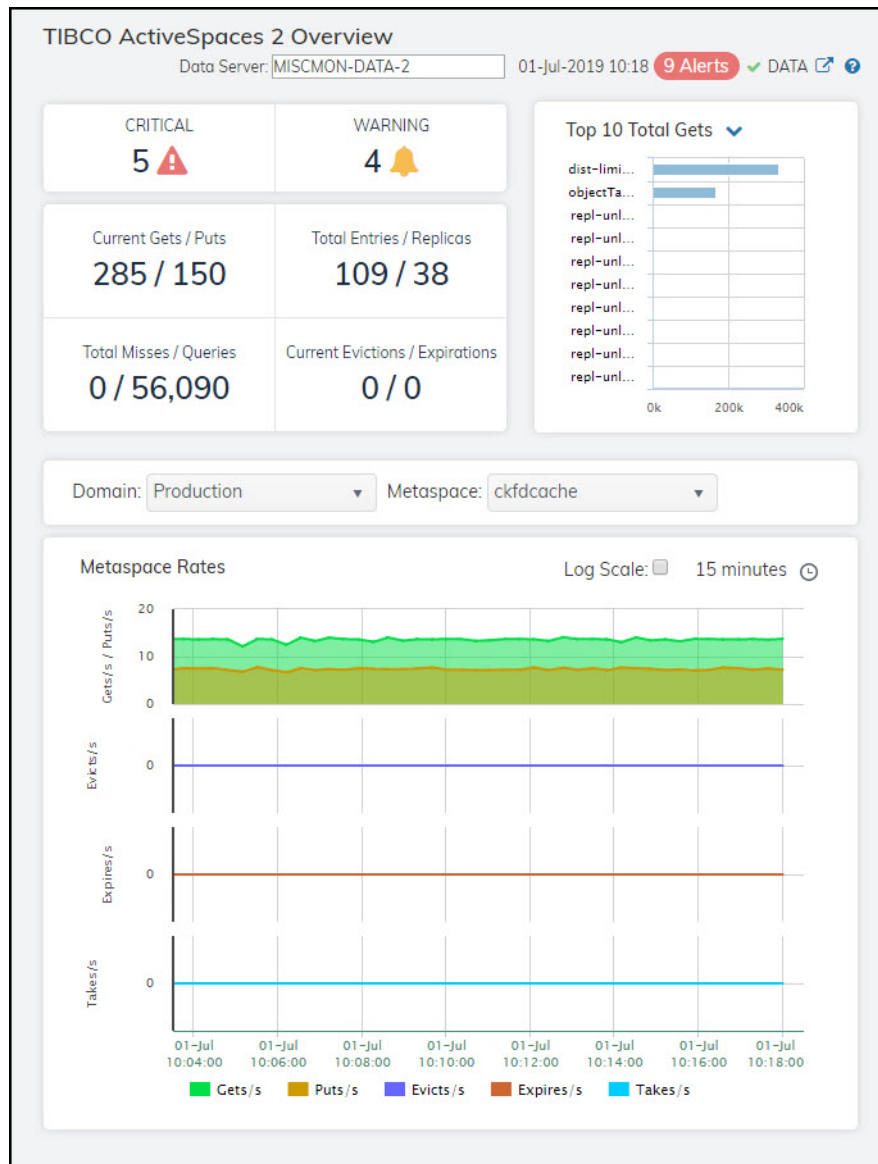
TIBCO ActiveSpaces 2 Overview

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”**: 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”**: 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”**: 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

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](#)” 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](#)” display, which shows metrics and trend data for a particular MetaSpace.

TIBCO ActiveSpaces 2 Metaspaces Table

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](#)” display.

Domain	Metaspace	Alert Level	Alert Count	Spaces	Members
Production	ckfdcache	⚠	6	41	3

Filter By:




Domain

Select the domain for which you want to view data.

Metaspaces

The total number of metaspaces found for the domain selected in the **Domain** dropdown, which are displayed in the **Metaspaces** table.

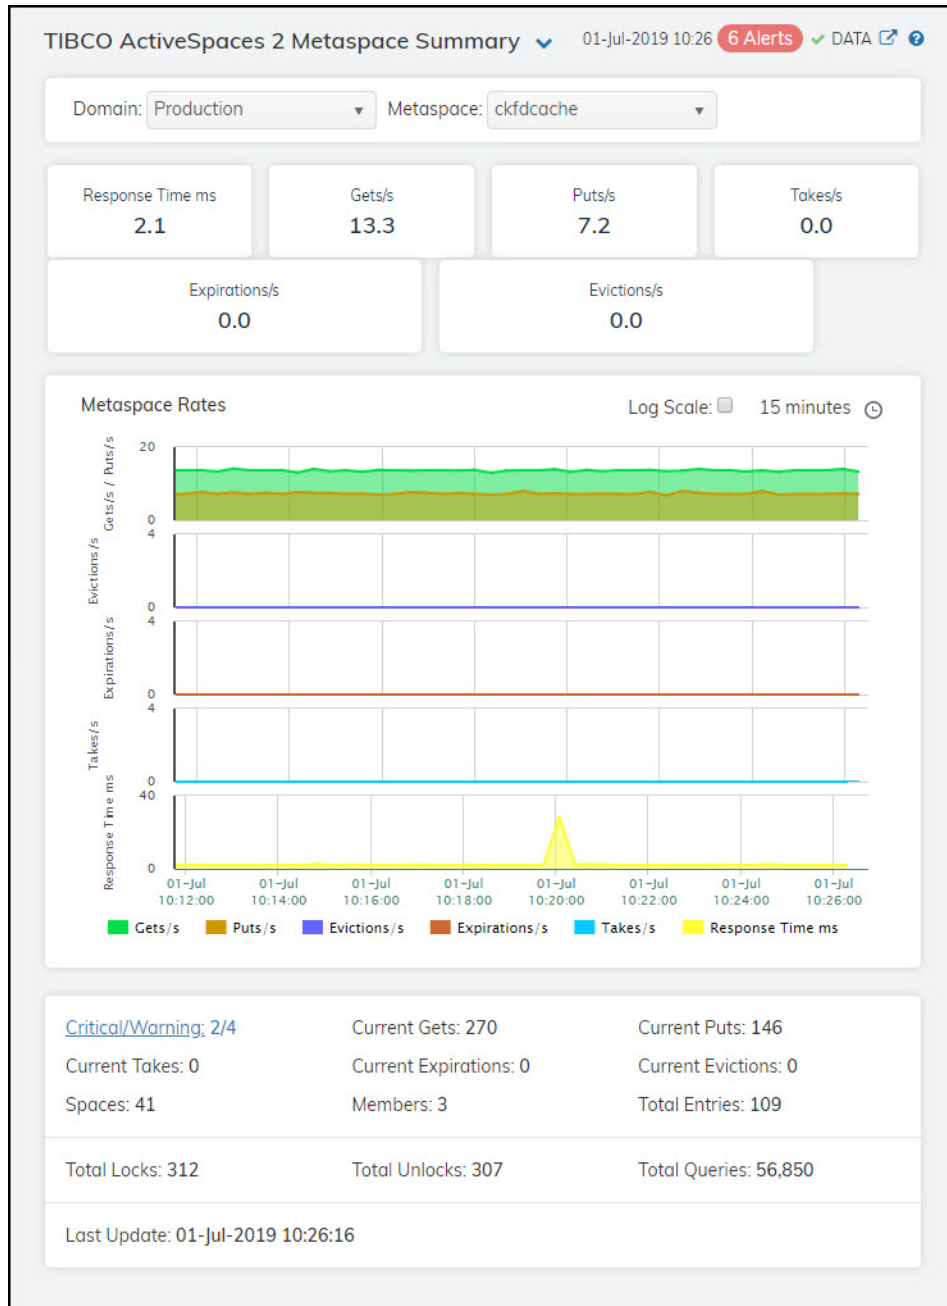
Metaspaces Table

Domain	The name of the domain.
Metaspace	The name of the metaspace.
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.
Spaces	The number of user spaces defined in the metaspace.*
Members	The number of members (clients and servers) associated with the metaspace.*
AS Version	The metaspace's current version of TIBCO ActiveSpaces.*
Total Entries	The total number of entries stored in the metaspace.*
Total Replicas	The total number of replicas stored in the metaspace.*
Response Time	The average response time for the metaspace.*
Total Gets	The total number of "get" operations performed on the user-spaces defined on the metaspace.*
Current Gets	The number of "get" operations performed on the user-spaces defined for the metaspace during the current polling interval.*
Gets/s	The rate of "get" operations (per second) performed on the user-spaces defined for the metaspace.*
Total Puts	The total number of "put" operations performed on the user-spaces defined on the metaspace.*
Current Puts	The number of "put" operations performed on the user-spaces defined for the metaspace during the current polling interval.*
Puts/s	The rate of "put" operations (per second) performed on the user-spaces defined for the metaspace.*
Total Takes	The total number of "take" operations performed on the user-spaces defined on the metaspace.*
Current Takes	The number of "take" operations performed on the user-spaces defined for the metaspace during the current polling interval.*
Takes/s	The rate of "take" operations (per second) performed on the user-spaces defined for the metaspace.*
Total Expirations	The total number of entries in the user-spaces defined on the metaspace that have expired.*
Current Expirations	The number of entries in the user-spaces defined for the metaspace that expired during the current polling interval.*
Expirations/s	The rate of entries in the user-spaces defined for the metaspace that expired (per second).*
Total Evictions	The total number of entries in the user-spaces defined on the metaspace that have been evicted.*
Current Evictions	The number of entries performed in the user-spaces defined for the metaspace that were evicted during the current polling interval.*

Evictions/s	The rate of entries in the user-spaces defined for the metaspace that were evicted (per second).*
Total Locks	The total number of locks in the user-spaces defined for the metaspace.*
Total Unlocks	The total number of unlocks in the user-spaces defined for the metaspace.*
Total Invokes	The remote invocation count.*
Total Queries	The browser queries count in the user-spaces defined for the metaspace.*
Total Misses	The total number of misses on the user-spaces defined for the metaspace.*
To Persist	The ToPersist count, which indicates how many tuples are required to be persisted to the database if the write-behind feature is configured.*
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.
Time Stamp	The date and time the row data was last updated.

TIBCO ActiveSpaces 2 MetaSpace Summary

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](#)" 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.



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:

- Response Time ms** The average response time for the metaspace. *
- Gets/s** The rate of "get" operations, per second.
- Puts/s** The rate of "put" operations, per second.
- Takes/s** The rate of "take" operations, per second.
- Expirations/s** The rate of expirations, per second.
- Evictions/s** The rate of evictions, per second.


Metaspace Rates Trend Graph

Traces the following:

- Gets/s** -- traces the number of gets per second.
- Puts/s** -- traces the number of puts per second.
- Evictions/s** -- traces the number of evicts per second.
- Expirations/s** -- traces the number of expirations per second.
- Takes/s** -- traces the number of takes per second.
- Response Time ms** -- traces the average response time 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.

Time Settings Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

Critical/Warning	The number of critical and warning alerts.
Current Takes	The total number of “take” operations for the metaspace for the current interval.
Spaces	The total number of spaces for the metaspace.
Current Gets	The total number of “get” operations for the metaspace for the current interval.
Current Expirations	The total number of expirations for the metaspace for the current interval.
Members	The total number of members for the metaspace.
Current Puts	The total number of “put” operations for the metaspace for the current interval.
Current Evictions	The total number of evictions for the metaspace for the current interval.
Total Entries	The total number of entries for the metaspace.
Total Locks	The total number of locks for the metaspace.
Total Unlocks	The total number of unlocks for the metaspace.
Total Queries	The total number of queries for the metaspace.
Last Update	The date and time in which the data in the display was last updated.

Spaces

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](#)” 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](#)”, which shows a heatmap view of all spaces contained within a particular metaspace.
- **Space Summary:** Opens the “[TIBCO ActiveSpaces 2 Space Summary](#)” display, which allows you to view metrics and trend data for a particular space.

TIBCO ActiveSpaces 2 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 double-clicking a row to view details for the selected space in the “[TIBCO ActiveSpaces 2 Space Summary](#)” 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

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 The number of spaces found in the query and listed in the table.




Table:

Domain The name of the domain.

Metaspace The name of the metaspace.

Space The name of the space.

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.

Space State The current state of the space.*

Members The total number of members in the space.*

Seeders The number of seeders in the space.*

Min Seeders The defined minimum seeder count (minimum number of seeders that need to be joined to the space before the space becomes ready).*

Capacity per Seeder The capacity value for the space in number of entries per seeder.*

Total Entries The total number of entries stored in the space.*

Total Replicas The total number of replicas stored in the space.*

Total Gets The total number of "get" operations performed on the user-spaces defined on the space.*

Current Gets The number of "get" operations performed on the user-spaces defined for the space during the current polling interval.*

Gets/s The rate of "get" operations (per second) performed on the user-spaces defined for the space.*

Total Puts The total number of "put" operations performed on the user-spaces defined on the space.*

Current Puts The number of "put" operations performed on the user-spaces defined for the space during the current polling interval.*

Puts/s The rate of "put" operations (per second) performed on the user-spaces defined for the space.*

Total Takes The total number of "take" operations performed on the user-spaces defined on the space.*

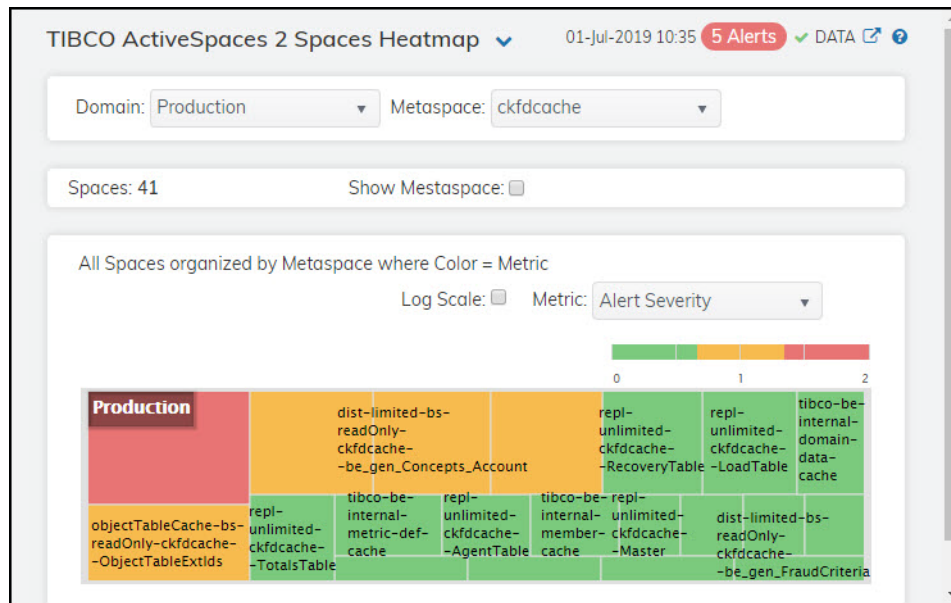
Current Takes The number of "take" operations performed on the user-spaces defined for the space during the current polling interval.*

Takes/s	The rate of "take" operations (per second) performed on the user-spaces defined for the space.*
Total Expirations	The total number of entries in the user-spaces defined on the space that have expired.*
Current Expirations	The number of entries in the user-spaces defined for the space that expired during the current polling interval.*
Expirations/s	The rate of entries in the user-spaces defined for the space that expired (per second).*
Total Evictions	The total number of entries in the user-spaces defined on the space that have been evicted.*
Current Evictions	The number of entries performed on the user-spaces defined for the space that were evicted during the current polling interval.*
Evictions/s	The rate of entries in the user-spaces defined for the space that were evicted (per second).*
Total Locks	The total number of locks in the user-spaces defined for the space.*
Total Unlocks	The total number of unlocks in the user-spaces defined for the space.*
Total Invokes	The remote invocation count.*
Total Queries	The total number of queries in the user-spaces defined for the space.*
Total Misses	The total number of misses in the user-spaces defined for the space.*
To Persist	The ToPersist count, which indicates how many tuples are required to be persisted to the database if the write-behind feature is configured.*
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.
Time Stamp	The date and time the row data was last updated.

TIBCO ActiveSpaces 2 Spaces Heatmap

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](#)" 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.

**Filter By:**

Domain Select the domain for which you want to see data.

Metaspace Select the metaspace for which you want to see data.

Spaces The total number of spaces found for the selected Domain/Metaspace combination.


Show Metaspace Select this check box to display the names of the metaspaces at the top of each rectangle in the heatmap.


Heatmap

Log 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.





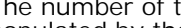
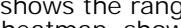

Metric Select the metric driving the heatmap display. The default is **Alert Severity**. Each **Metric** has a color gradient bar that maps values to colors. The heatmap is organized by spaces, where each rectangle represents a space. Mouse-over any rectangle to display the current values of the metrics for the space. Click on a rectangle to drill-down to the associated ["TIBCO ActiveSpaces 2 Space Summary"](#) 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.

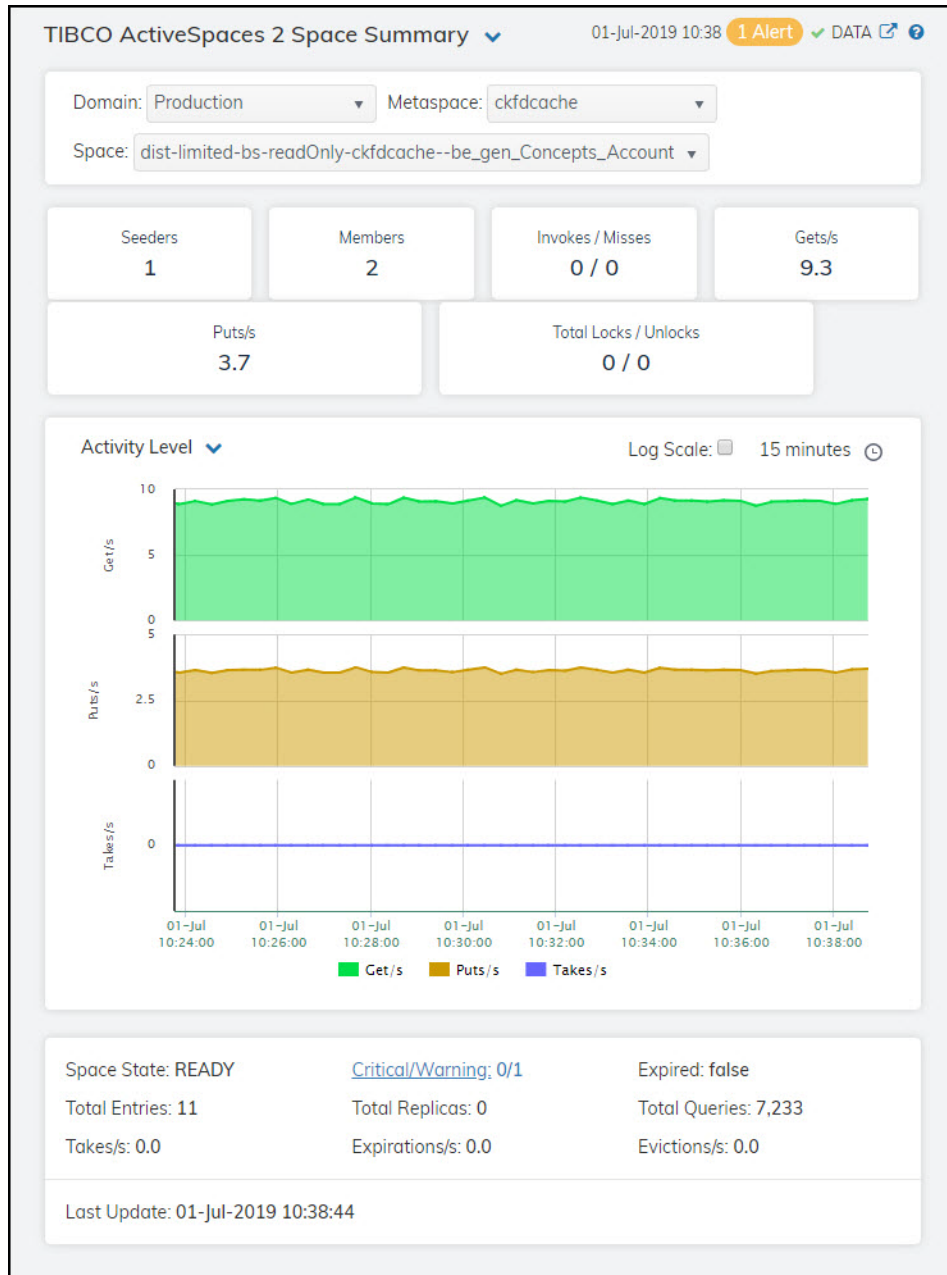
 Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.

 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.
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. 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. 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.

TIBCO ActiveSpaces 2 Space Summary

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](#)", 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.



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:

Seeders	The number of seeders in the space.*
Members	The total number of members associated with the space.*
Invokes/ Misses	The total number of invocations and the total number of misses on the user-spaces defined for the space.*
Gets/s	The number of "get" operations received per second.
Puts/s	The number of "put" operations received per second.
Total Locks/ Unlocks	The total number of locks and unlocks in the user-spaces defined for the space.*

Trend Graphs**Activity Level**

Gets/s -- traces the number of gets per second.

Puts/s -- traces the number of puts per second.

Takes/s -- traces the number of takes per second.

Exceptions

Expires/s -- traces the number of expires per second.

Evicts/s -- traces the number of evicts per second.


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.

Time Settings Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

Space State	The current state of the space.*
Total Entries	The total number of entries stored in the space.*
Takes/s	The number of takes received per second.
Critical/Warning	The number of critical and warning alerts.
Total Replicas	The total number of replicated entries in the space.*
Expirations/s	The number of expirations received per second.
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.
Total Queries	The total number of queries in the user-spaces defined for the space.*
Evictions/s	The number of evictions received per second.
Last Update	The date and time in which the data in the display was last updated.

Members

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”](#) 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”](#), which shows a heatmap view of all members in a particular metaspace.
- **Single Member**: Opens the [“TIBCO ActiveSpaces 2 Member Summary”](#) display, which allows you to view current and trending data for a single member for a particular metaspace.

TIBCO ActiveSpaces 2 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 double-clicking a row to view details for the selected member in the [“TIBCO ActiveSpaces 2 Member Summary”](#) display.

Domain	Metaspace	Member Name	Alert Level	Alert Count	Management Role
Production	ckfdcach	7f000001-e48e	✓	0	MEMBER
Production	ckfdcach	7f000001-e4f2	✓	0	MEMBER
Production	ckfdcach	7f000001-e556	✓	0	MANAGER

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.

Members The resulting total number of members found in the filtered query, and listed in the **Members** table.

Table

Domain The name of the domain.

Metaspace The name of the metaspace.

Member Name The name of the member.

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.

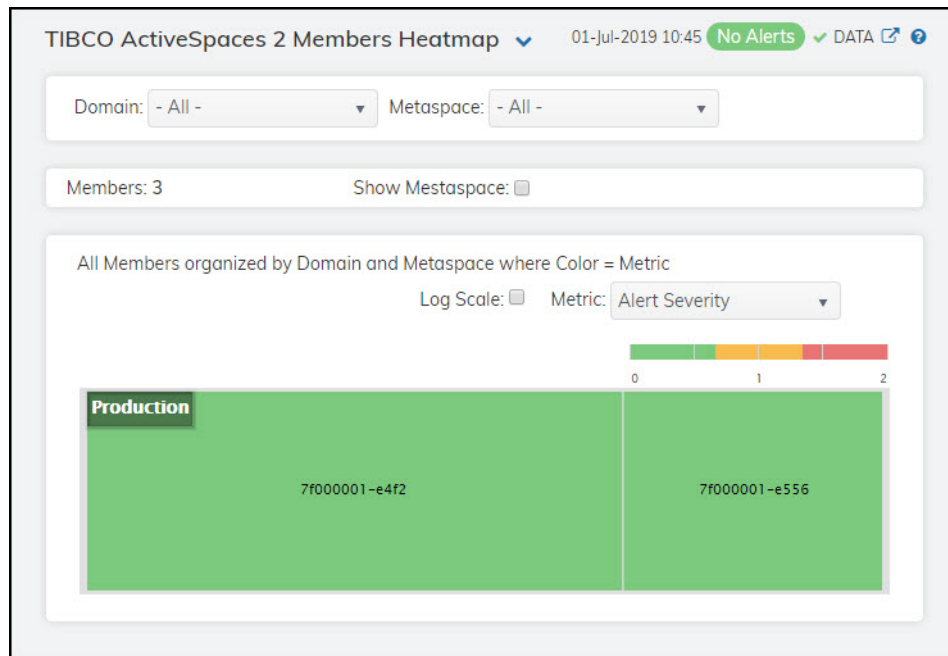
Management Role	The member's role within the metaspace.
Host Address	The IP address of the host.*
Host Port	The port of the host.*
PID	The process ID of the process being monitored.*
Process Name	The name of the process.*
Spaces	The number of spaces in which the metaspace member is a member.*
Total Entries	The number of entries associated with the member.*
Total Replicas	The number of replicas.*
Total Gets	The total number of "get" operations performed on the user-spaces defined on the metaspace.*
Current Gets	The number of "get" operations performed on the user-spaces defined for the metaspace during the current polling interval.*
Gets/s	The rate of "get" operations (per second) performed on the user-spaces defined for the metaspace.*
Total Puts	The total number of "put" operations performed on the user-spaces defined on the metaspace.*
Current Puts	The number of "put" operations performed on the user-spaces defined for the metaspace during the current polling interval.*
Puts/s	The rate of "put" operations (per second) performed on the user-spaces defined for the metaspace.*
Total Takes	The total number of "take" operations performed on the user-spaces defined on the metaspace.*
Current Takes	The number of "take" operations performed on the user-spaces defined for the metaspace during the current polling interval.*
Takes/s	The rate of "take" operations (per second) performed on the user-spaces defined for the metaspace.*
Total Expirations	The total number of entries in the user-spaces defined on the metaspace that have expired.*
Current Expirations	The number of entries performed in the user-spaces defined for the metaspace that expired during the current polling interval.*
Expirations/s	The rate of entries in the user-spaces defined for the metaspace that expired (per second).*
Total Evictions	The total number of entries in the user-spaces defined on the metaspace that have been evicted.*
Current Evictions	The number of entries performed in the user-spaces defined for the metaspace that were evicted during the current polling interval.*
Evictions/s	The rate of entries in the user-spaces defined for the metaspace that were evicted (per second).*
Total Locks	The total number of locks in the user-spaces defined for the metaspace.*
Total Unlocks	The total number of unlocks in the user-spaces defined for the metaspace.*
Total Invokes	The remote invocation count.*

Total Queries	The total number of queries in the user-spaces defined for the metaspace.*
Total Misses	The total number of misses in the user-spaces defined for the metaspace.*
To Persist	The ToPersist count, which indicates how many tuples are required to be persisted to the database if the write-behind feature is configured.*
Join Time	The time that the member joined the space.
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.
Time Stamp	The date and time the row data was last updated.

TIBCO ActiveSpaces 2 Members Heatmap

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](#)" 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 Metaspace** 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.

**Filter By:**

Domain Select the Domain for which you want to view data.

Metaspace Select the metaspace for which you want to view data.

Members The number of members found for the selected **Domain/Metaspace** combination.

Show Metaspace Select this check box to display the names of the metaspaces at the top of each rectangle in the heatmap.

Heatmap

Log 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.


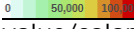
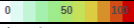
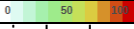
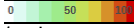
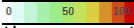
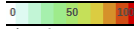
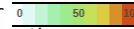
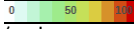
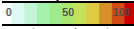
Metric Select the metric driving the heatmap display. The default is **Alert Severity**. Each **Metric** has a color gradient bar that maps values to colors. The heatmap is organized by members, where each rectangle represents a member. Mouse-over any rectangle to display the current values of the metrics for the member. Click on a rectangle to drill-down to the associated ["TIBCO ActiveSpaces 2 Member Summary"](#) 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.

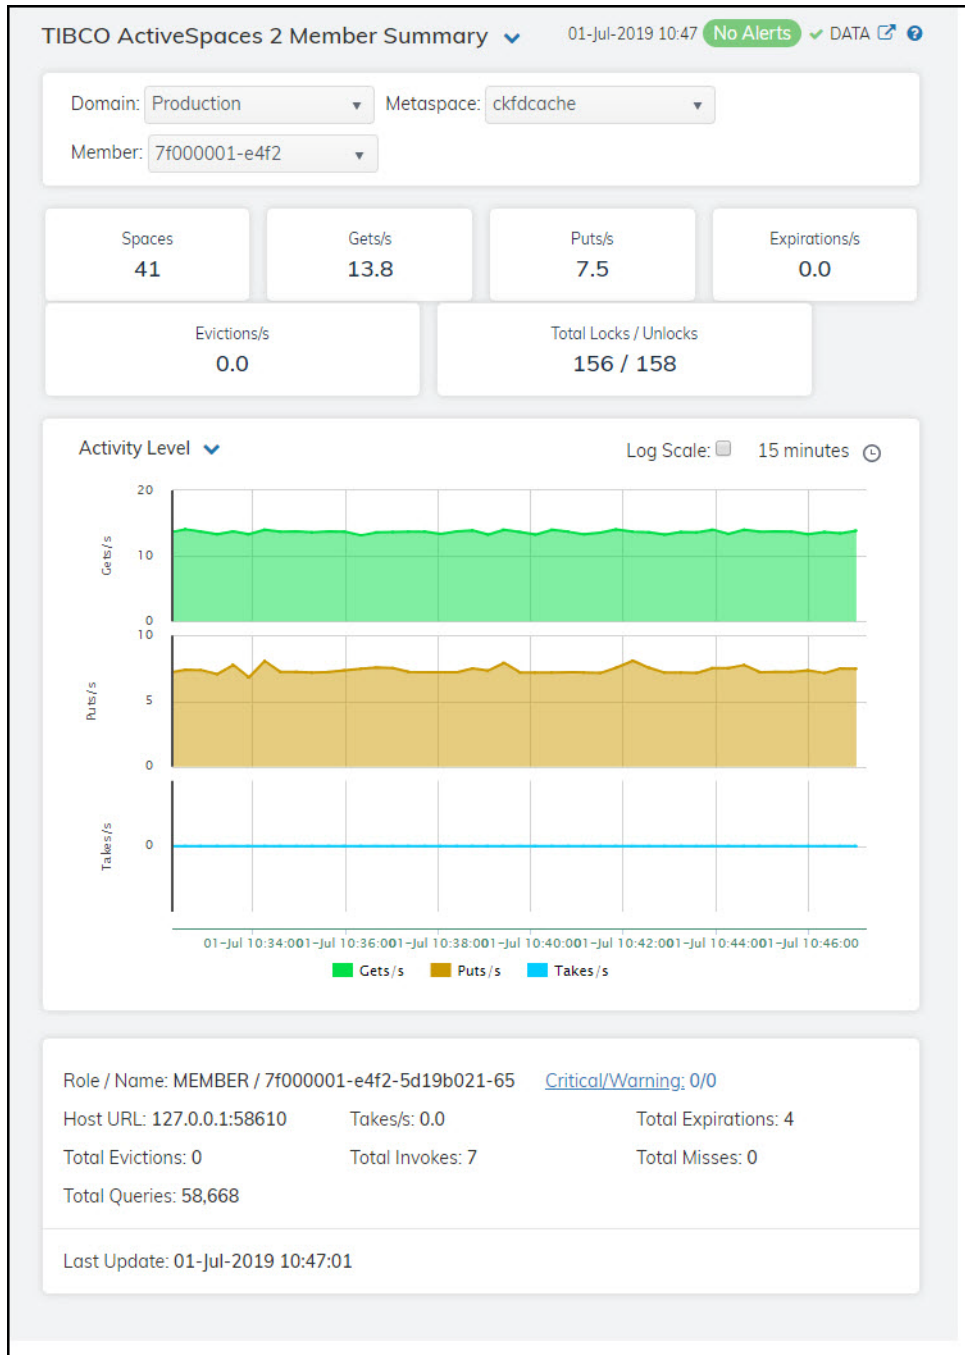
Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.

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

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"](#) 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.



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:

- Spaces** The total number of spaces in which the member is a member.*
- Gets/s** The rate of "get" operations (per second) performed on the user-spaces defined for the metaspace.*
- Puts/s** The rate of "put" operations (per second) performed on the user-spaces defined for the metaspace.*
- Expirations/s** The rate of entries in the user-spaces defined for the metaspace that expired (per second).*
- Evictions/s** The rate of entries in the user-spaces defined for the metaspace that were evicted (per second).*
- Total Locks/
Unlocks** The total number of locks and unlocks in the user-spaces defined for the metaspace.*

Trend Graphs**Activity Level**

- Gets/s** -- traces the number of gets per second.
- Puts/s**-- traces the number of puts per second.
- Takes/s** -- traces the number of takes per second.


Exceptions

- Expirations/s** -- traces the total number of expires, or the number of expires per second with **Use Rates** selected.
- Evictions/s** -- traces the total number of evicts, or the number of evicts per second with **Use Rates** selected.

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 Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

Role/Name	The role of the member and the name of the member.
Host URL	The IP address of the host.
Total Evictions	The total number of entries in the user-spaces defined on the metaspace that have been evicted.*
Total Queries	The total number of queries in the user-spaces defined for the metaspace.*
Critical/Warning	The number of critical and warning alerts.
Takes/s	The rate of “take” operations (per second) performed on the user-spaces defined for the metaspace.*
Total Invokes	The remote invocation count.*
Total Expirations	The total number of entries in the user-spaces defined on the metaspace that have expired.*
Total Misses	The total number of misses in the user-spaces defined for the metaspace.*
Last Update	The date and time in which the data in the display was last updated.

TIBCO Adapters

The following Views and their associated displays are available in the Monitor. This section describes the Monitor displays and includes:

- [“TIBCO Adapters Overview”](#): Describes the **TIBCO Adapters Overview** display.
- [“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.

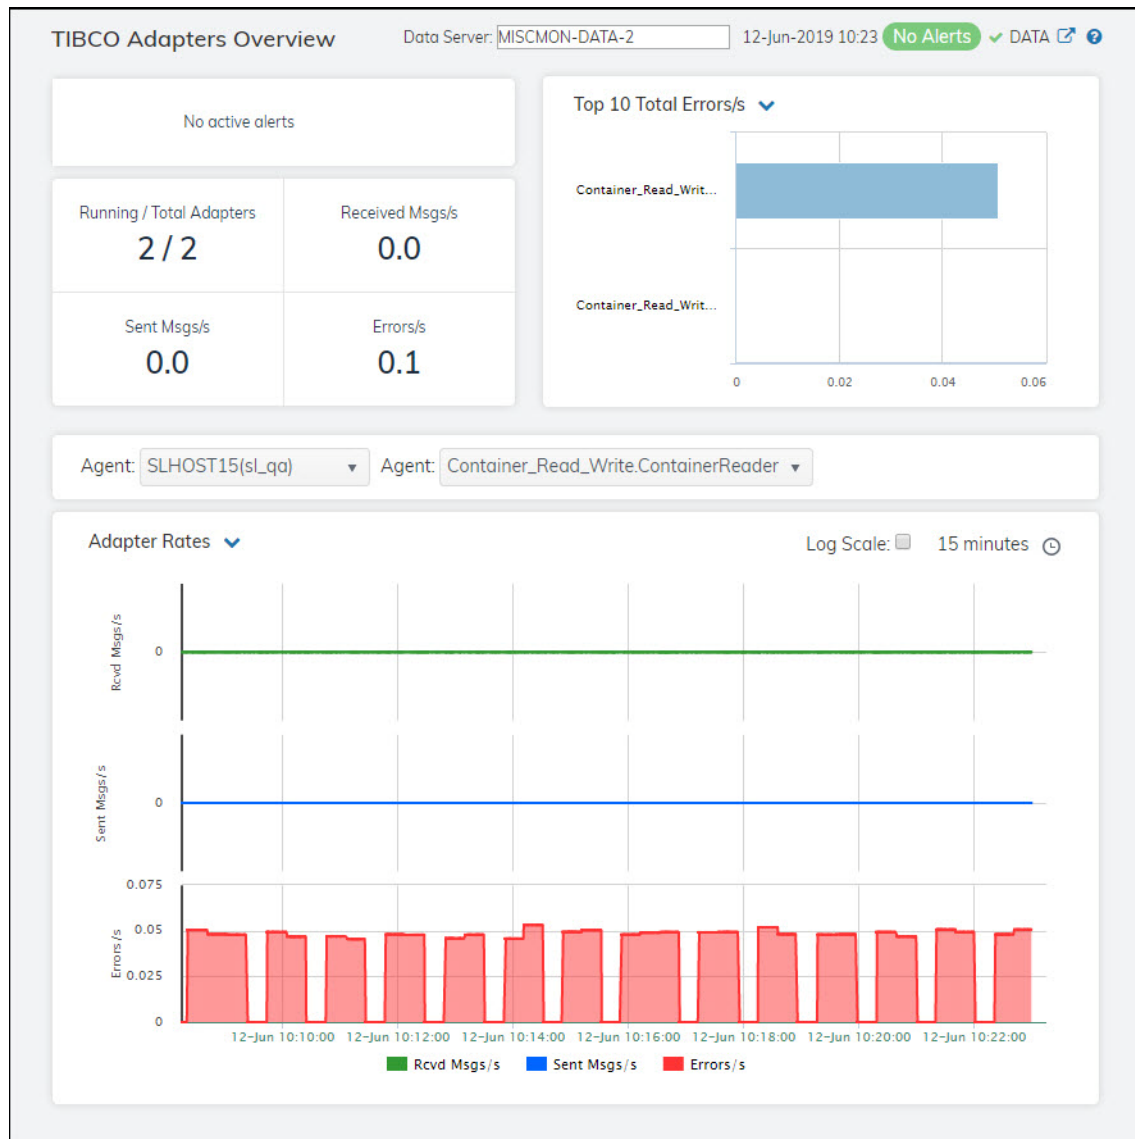
TIBCO Adapters Overview

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

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](#)” 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](#)” display, which provides a view of all adapters in a heatmap format and their associated metrics.
- **Single Adapter**: Opens the “[TIBCO Adapter Summary](#)” display, which provides current and historical data for a single adapter.

TIBCO 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 double-clicking a row to view details for the selected adapter in the [“TIBCO Adapter Summary”](#) 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

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:

Adapters The total number of adapters listed in the table.

All Adapters Table:

Agent Name The name of the agent.

Application Instance The name of the application instance.

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.

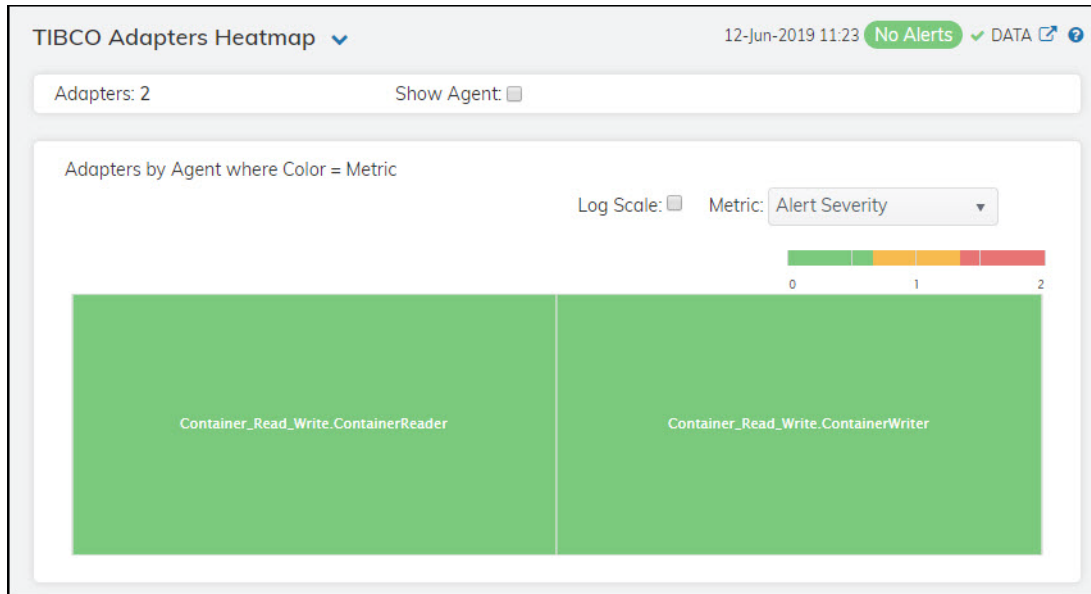
Adapter Name	The name of the adapter.*
Received Msgs	The number messages received.*
Sent Msgs	The total number of messages sent.*
Total Errors	The total number of errors.*
Total Errors/s	The number of errors per second.*
Sent Msgs/s	The number of messages sent per second.*
New Errors	The number of new errors received since the last polling update.*
Process ID	The process ID of TIBCO Adapter you are running.*
Current Sent Msgs	The increase in the number of messages sent (from the previous polling period to the current polling period).*
Last Restart	The date and time the adapter was last restarted.*
Current Received Msgs	The increase in the number of messages received (from the previous polling period to the current polling period).*
Current Total Errors	The increase in the number total errors (from the previous polling period to the current polling period).*
Received Msgs/s	The number of messages received per second.*
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 Adapters > 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.

TIBCO Adapters Heatmap

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.


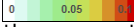
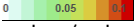
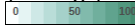
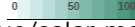
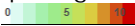
Each rectangle in the heatmap represents an adapter. The rectangle color indicates the most critical alert state associated with the adapter. Choose a different metric to display from the **Metric** drop-down menu. Mouse over a rectangle to see additional metrics. By default, this display shows **Alert 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”](#) display.



Fields and Data:

- Adapters** The number of adapters included in the display.
- Show Agent** Select this check box to display the names of the agents 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.
- 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 Scale** is not selected.
- 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 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”](#) 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	<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 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.</p>
Sent Msgs/s	<p>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 TadAdapterMsgsSentRateHigh. The middle value in the gradient bar indicates the middle value of the range.</p> <p>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.</p>
Received Msgs/s	<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 TadAdapterMsgsRcvdRateHigh. The middle value in the gradient bar indicates the middle value of the range.</p> <p>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.</p>
Current Sent Msgs	<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 Auto Scale check box does not impact this metric.</p>
Current Received Msgs	<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 Auto Scale check box does not impact this metric.</p>
Current Errors	<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 TadAdapterDeltaErrorsHigh. The middle value in the gradient bar indicates the middle value of the range.</p> <p>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.</p>

TIBCO Adapter Summary

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](#)" 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 🕒

Uptime: 9 seconds
[Critical/Warning: 0/0](#)
New Errors: 1

Last Update: 12-Jun-2019 11:37:31

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:

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:

Received Msgs/s	The number of messages received per second on the adapter.*
Sent Msgs/s	The number of messages sent per second from the adapter.*
Errors/s	The number of errors occurring per second on the adapter.*
Current Received Msgs	The increase in the number of messages received since the last polling update.*
Current Sent Msgs	Deltas -- The increase in the number of messages sent since the last polling update.*
Current Errors	The increase in the number of errors since the last polling update.*

Adapter Service Information Table

Service Name	The name of the service.*
Subject	The name of the subject.*
Type	The type of adapter service.*
# Messages	The current number of messages.*

Host Information Table

Name	The name of the host.*
Value	The host's value.*


Trend Graphs**Adapter Rates**

- Rcvd Msgs/s** -- traces rate of messages received.
- Sent msgs/s**-- traces the rate of messages sent.
- Errors/s** -- traces the rate of errors.

Adapter Current Counts

- Rcvd Msgs** -- traces the increase in the number of messages received since the last polling update.
- Sent Msgs** -- traces the increase in the number of messages sent since the last polling update.
- Errors** -- traces the increase in the number of errors since the last polling update.

- 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** Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

- Uptime** The amount of time the adapter has been up and running.
- Critical/Warning** The number of critical and warning errors.
- New Errors** The number of errors since the last data update.
- Last Update** The date and time in which the data in the display was last updated.

TIBCO BusinessEvents

The HTML version features an overview display, "[TIBCO BusinessEvents Overview](#)" (pictured below), and the following Views which can be found under **Components** tab > **Middleware** > **TIBCO BusinessEvents**.

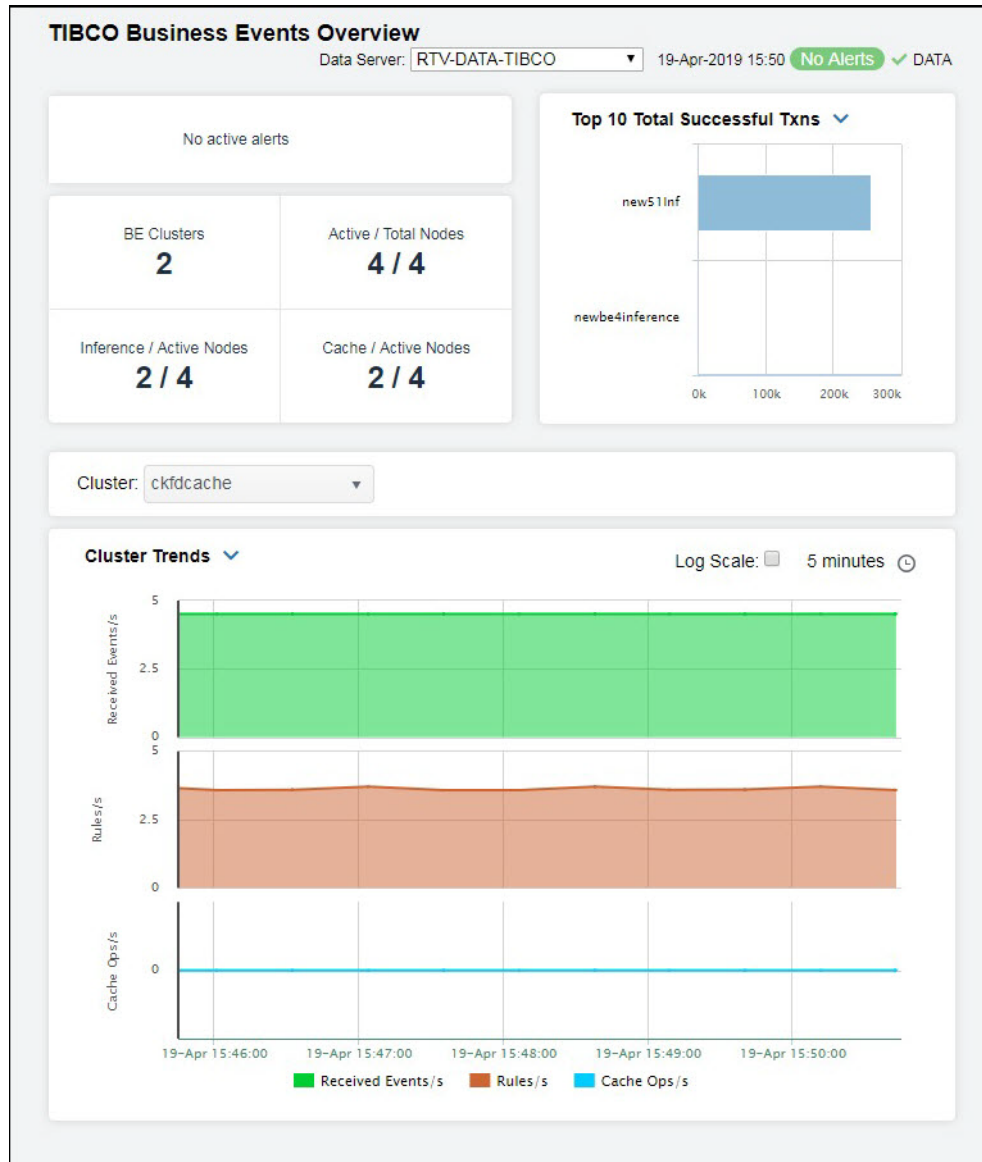
TIBCO BusinessEvents Overview

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.



The TIBCO BusinessEvents Monitor contains the following Views:

- "BE Clusters"
- "BE Nodes"
- "BE Events"
- "BE Concepts"

BE Clusters

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](#)” 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](#)”, which shows cluster and alert status for all BE clusters.
- **Clusters Summary**: Opens the “[TIBCO BE Cluster Summary](#)” display, which shows information for a single BE Cluster.

TIBCO BE Clusters Table




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](#)” 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.

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	

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.

Clusters:	The total number of clusters in the table.
Cluster Name	The name of the TIBCO BusinessEvents cluster.
Alert Level	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.
Alert Count	The total number of critical and warning alerts.
Member Count	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. Note: 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.
Events Received	The total number of events received. *
Events Sent	The total number of events sent. *
Received Events/ S	The rate of events received in the cluster.
Asserts/s From Channel	The rate of events asserted into the Rete network via the channel.
Retracted/s From Channel	The rate of events retracted/deleted from the Rete network via the channel.
Modified/s From Channel	The rate of events modified in the Rete network via the channel.
Fired Rules/s	The rate of rules fired in the cluster.
Concept Max Get Avg Time ms	The longest time taken for a "get" operation for any node in the cluster since the cluster was started (in milliseconds). *
Max Concept Put Avg Time ms	The longest time taken for a "put" operation for any node in the cluster since the cluster was started (in milliseconds). *
Max Concept Remove Avg Time ms	The longest time taken for a "remove" operation for any node in the cluster since the cluster was started (in milliseconds). *
Max Concept Op Time	The longest time taken for a concept operation (get/put/remove) for any node in the cluster since the cluster was started. *
Concept Gets/s	The rate of "get" operations in the cluster.

Concept Puts/s	The rate of "put" operations in the cluster.
Concept Removes/s	The rate of "remove" operations in the cluster.
Concept Ops/s	The rate of operations (gets/puts/removes) in the cluster.
Backing Store Max Erase Time	The longest time taken for an "erase" operation in the Backing Store for any node in the cluster.*
Backing Store Max Load Time	The longest time taken for a "load" operation in the Backing Store for any node in the cluster.*
Backing Store Max Store Time	The longest time taken for a "store" operation in the Backing Store for any node in the cluster.*
Backing Store Max Op Time	The longest time taken to perform an operation (erase/load/store) in the Backing Store for any node in the cluster.*
Backing Store Erases/s	The rate of "erases" in the Backing Store.
Backing Store Loads/s	The rate of "loads" into the Backing Store.
Backing Store Stores/s	The rate of "stores" into the Backing Store.
Backing Store Ops/s	The rate of operations (erases/loads/stores) in the Backing Store.
Source	The name of the data server from which the data was collected.
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.
Time Stamp	The date and time, relative to the Data Server, that data was last collected for the engine.

TIBCO BE Clusters Heatmap

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](#)" 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.










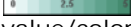
Fields and Data

Clusters The number of clusters shown in the heatmap.

All Clusters Heatmap

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.

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"](#) display for a detailed view of metrics for that particular cluster.

Alert Severity	<p>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:</p> <ul style="list-style-type: none">  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	<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 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.</p>
Member Count	<p>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.</p>
Events Received	<p>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.</p>
Events Sent	<p>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.</p>

TIBCO BE Cluster Summary

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"](#) display or the ["TIBCO BE Events Table"](#) 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.



Filter By:

The display might include these filtering options:

Cluster Choose a cluster for which you want to see metrics.

Fields and Data

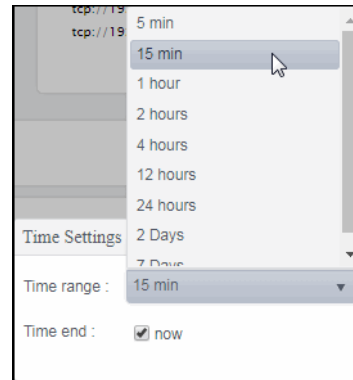
This display includes:

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.

Active Nodes	Lists the number of active nodes on the cluster.
Rules Fired/s	The rate of rules fired in the Rete network.
Received Events/s	The number of events received since the last data update. *
Gets/s / Puts/s	The rate of "get" operations in the L1 cache, and the rate of "put" operations in the cache.
Removes/s	The rate of "removes" in the cache.
Total Cache Ops/s	The rate of operations (gets/puts/removes) in the cache.
Trend Graphs	<p>Cluster Trends Shows the following metrics for the selected cluster.</p> <ul style="list-style-type: none"> Received Events/s -- Traces the rate of events received in the cluster. Rules/ sec -- Traces the rate of rules in the cluster. Cache Ops/ sec -- Traces the rate of cache operations in the cluster. <p>Backing Store Trends Shows the following metrics for the selected cluster.</p> <ul style="list-style-type: none"> Backing Store Ops/sec-- Traces the rate of backing store operations in the cluster. Max Backing Store Avg Time-- Traces the average maximum time per backing store operation. <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.</p>

Time Settings

Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

Nodes

Lists the number of nodes on the cluster.

Concept Max Get Avg Time ms

The longest time taken for a "get" operation for any node in the cluster since the cluster was started (in milliseconds).*

Max Concept Remove Avg Time ms

The longest time taken for a "put" operation for any node in the cluster since the cluster was started.*

Critical/Warning

The number of critical and warning alerts on the cluster.

Max Concept Put Avg Time ms

The longest time taken for a "put" operation for any node in the cluster since the cluster was started.*

Cache Nodes

The number of cache nodes on the cluster.

Last Update

The date and time the data was last updated in the display.

BE Nodes

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](#)” 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](#)”, which shows cluster and alert status for all BE cluster nodes.
- **Inference Node Summary**: Opens the “[TIBCO BE Inference Node Summary](#)” display, which shows information for (inference) agents for a single BE cluster node.
- **Storage Node Summary**: Opens the “[TIBCO BE Storage Node Summary](#)” display, which displays cache data for a specific node.

TIBCO BE Cluster Nodes Table

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.

TIBCO BE Cluster Nodes Table 19-Apr-2019 16:06 No Alerts DATA

Cluster: - All -

Nodes: 4

Cluster Name	Connection	Alert Level	Alert Count	MemberCount	Auto S
ckfdcache	new51Cache	✓	0	2	▼
ckfdcache	new51Inf	✓	0	2	▼
fdcache	newbe4cache	✓	0	2	▼
fdcache	newbe4inference	✓	0	2	▼

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.

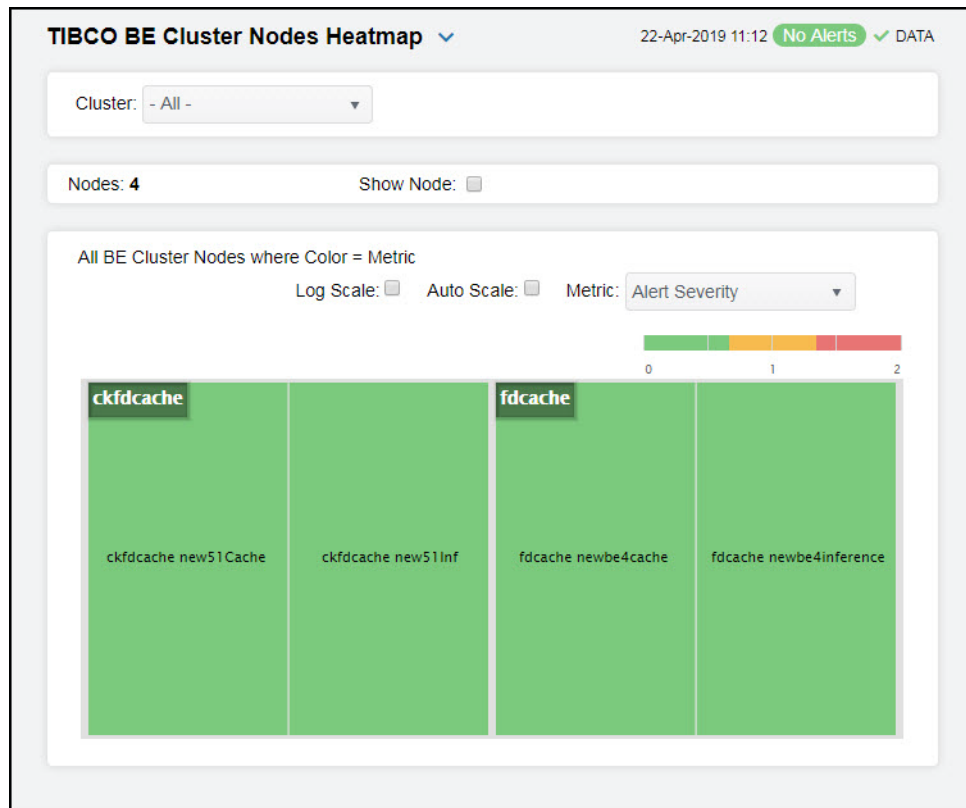
Nodes:	The total number of clusters in the table.
Cluster Name	The name of the TIBCO BusinessEvents cluster.
Connection	The name of the node.
Alert Level	The severity level of open alerts. Values range from 0 to 2 , where 2 is the greatest Severity:  One or more alerts exceeded their ALARM LEVEL threshold.  One or more alerts exceeded their WARNING LEVEL threshold.  No alert thresholds have been exceeded.
Alert Count	The total number of critical and warning alerts.
Member Count	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 Clusters table. If they do not match, the node did not join the cluster properly and, hence, the cluster should be restarted.
Auto Startup	When checked (true), this feature is enabled.
Backing Store Enabled	When checked (true), this feature is enabled.*
Cache Aside	When checked (true), this feature is enabled.*
Serialization Optimized	When checked (true), this feature is enabled.*
Storage Enabled	When checked (true), this feature is enabled.*
Cache Type	The type of TIBCO BusinessEvents cache.*
BE Version	The approximate TIBCO BusinessEvents version, as configured by the connection property. The exact version information is not available via JMX.
isCacheNode?	When checked (true), 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.
Node ID	A unique string that identifies the node.
CPU %	The amount of CPU, in percent, used by the node. This value is derived from the java.lang.OperatingSystem MBean.
Max Heap	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.

Used Heap	The current amount of memory, in megabytes, in use by the JVM for heap space. This value is provided by standard Java MBeans.
NonHeapMemoryUsage.max	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.
NonHeapMemoryUsage.Used	The current amount of memory, in megabytes, in use by the process (not counting heap usage). This value is provided by standard Java MBeans.
Connected	When checked (true), the node is currently connected to the Data Server via JMX.
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.
Time Stamp	The date and time, relative to the Data Server, that data was last collected for the node.

TIBCO BE Cluster Nodes Heatmap





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.

**Filter By:**

The display might include these filtering options:

- | | |
|--------------------------|--|
| Cluster | Choose a cluster for which you want to see metrics. |
| Nodes: | The total number of nodes in the display. |
| Show Node | Select this check box to display the name of the node in each rectangle in the heatmap. |
| 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. |
| Auto Scale | 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). |
| 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. |

Alert Severity	<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 2 is the greatest Alert Severity.</p> <p>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.</p> <p>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.</p> <p>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.</p>
Alert Count	<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 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>
JVM % CPU Used	<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 0 to the alert threshold of JvmCpuPercentHigh. The middle value in the gradient bar indicates the middle value of the range.</p>
JVM % Heap Used	<p>The total percentage of JVM Heap 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 0 to the alert threshold of JvmMemoryUsedHigh. The middle value in the gradient bar indicates the middle value of the range.</p>

TIBCO BE Inference Node Summary

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"](#) display, the ["TIBCO BE Agent Event Summary"](#) 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 🕒

■ CPU % ■ Max Heap ■ Used Heap

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**

Filter By:

The display might include these filtering options:

Node	Select a node for which you want to view metrics.
Agent ID	Select the agent ID for which you want to view metrics.

Fields and Data:

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.

Queue Capacity %	The queue capacity for the agent.*
Top Rules/s	The maximum rules fired per second for all inference agents from the selected node.
Top Job Rate	The maximum number of jobs per second across all inference agents from the selected node.
Top Hit Ratio	The maximum hit ratio across all inference agents from the selected node.
Heap Used %	The percent heap utilization of this inference node.
CPU %	The percentage of CPU utilization on this inference node.

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.

Agent ID	The agent's ID.
Agent Class	The agent's class. See TIBCO documentation for more information.
Type	The type of agent (Inference, Cache, Query, or Dashboard).*
Current State	The current state of the agent.*
Started	When checked, denotes that the agent is started.*
Suspended	When checked, denotes that the agent is suspended.*
Concurrent	When checked, denotes that it is a concurrent agent.*
Queue Capacity	The queue capacity for the agent.*
Queue Size	The queue size for the agent.*
Thread Count	The total number of threads for the agent.*
Total # Rules Fired	The total number of rules fired for the agent.*
Rules Fired	The number of rules fired.*
Rules/s	The rate of rules fired for the agent.

Avg Receive Time	See TIBCO documentation for more information.*
Avg Txn Commit Time	The average amount of time taken to commit a transaction.*
Cache Queue Remaining	The total amount of remaining space on the cache queue.*
DB Ops Queue Remaining	The total amount of remaining space on the DB Operations queue.*
Hit Ratio	See TIBCO documentation for more information.*
Job Rate	See TIBCO documentation for more information.*
L1 Cache Max Size	The maximum size of the L1 cache.*
L1 Cache Size	The current size of the L1 cache.*
Max Active	See TIBCO documentation for more information.*
# Event Threads	The total number of currently active event threads.*
# Jobs	The total number of currently active jobs.*
Priority	See TIBCO documentation for more information.*
Read Only	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.
Time Stamp	The date and time, relative to the Data Server, that data was last collected for the agent.

Trend Graph**Utilization**

Shows metrics for the selected node.

% CPU -- Traces the amount of CPU used, in percent, by the node.

Max Heap -- Traces the maximum amount of heap space, in bytes, used by the node since the agent was started.

Used Heap -- Traces the current amount of heap space, in bytes, used by the agent.

Rules/s and Threads

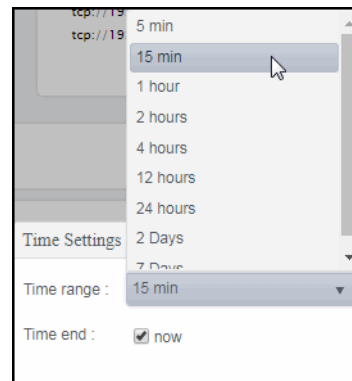
Shows metrics for the selected node.

Rules/s -- Traces the number of rules processed, per second, by the agent.

Thread Count-- Traces the number of threads being used by the node.

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.

Time Settings Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

Connection	The JMX connection method specified in the connection property for a given engine. It is displayed as either a combination of the host and port fields (< host >:< port >), 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.
Max Number of Jobs	The maximum number of jobs.*
Max Number of Event Threads	The maximum number of event threads.*
Critical/Warning	The number of critical and warning alerts.

Max Number of Rules Fired	The maximum rules fired for all inference agents from the selected node.
Max Avg Receive Time ms	The maximum average receive time for all inference agents from the selected node.
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.
Max Threads	The maximum number of threads on the selected node.
Max Avg Commit Time ms	The maximum average commit time, in milliseconds, for all inference agents on the selected node.
Last Update	The date and time the data in the display was last updated.

TIBCO BE Storage Node Summary

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**

Filter By:

The display might include these filtering options:

Node: Select the node for which you want to see metrics.

Fields and Data

CPU %	The amount of CPU, in percent, used by the node. This value is provided by standard Java MBeans.
Top Load Avg Time ms	The maximum load average time, in milliseconds, across all backing stores of this storage node.
Top Store Avg Time ms	The maximum average store time, in milliseconds, across all backing stores of this storage node.
Heap Used %	The percentage of heap utilization of this storage node.
Top Loads/s	The maximum number of loads per second across all backing stores of this storage node.
Top Stores/s	The maximum number of stores per second across all backing stores of this storage node.

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.

Cache Name	The name of the cache. *
Active	When checked, denotes that the cache is active. *
Delete Avg Time	The average amount of time taken for a "delete" ("erase") operation in the Backing Store for the cache. *
Load Avg Time	The average amount of time taken for a "load" operation in the Backing Store for the cache. *
Store Avg Time	The average amount of time taken for a "store" operation in the Backing Store for the cache. *
Delete Total	The total number of "delete" operations in the Backing Store for the cache. *
Load Total	The total number of "load" operations in the Backing Store for the cache. *
Store Total	The total number of "store" operations in the Backing Store for the cache. *
Deletes	The number of "delete" operations during the last polling interval. *
Loads	The number of "load" operations during the last polling interval. *
Stores	The number of "store" operations during the last polling interval. *
Deletes/s	The rate of "delete" operations in the node.
Loads/s	The rate of "load" operations in the node.
Stores/s	The rate of "store" operations in the node.

Expired

When checked (true), 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.

Trend Graphs**Utilization**

Shows metrics for the selected cluster/node combination:

% CPU -- Traces the amount of CPU used, in percent, by the engine.

Max Heap (MB)-- Traces the maximum amount of memory, in megabytes, that can be used by the JVM for heap space.

Used Heap (MB) -- Traces the used heap space, in megabytes, in use by the JVM.

Table Sizes

Shows metrics for the selected cluster/node combination:

Table Size -- Traces the number of unique objects cached in the local index table.


Ext ID Tbl Size -- Traces the number of entries in the table of external IDs used as indexes by the backing store.

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.

Time Settings

Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

Connection

The JMX connection method specified in the **connection** property for a given engine. It is displayed as either a combination of the host and port fields (**<host>:<port>**), 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.

Expired

When **true**, 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.

Storage Enabled

When checked (**true**), this feature is enabled.

Critical/Warning

The number of critical and warning alerts.

Node ID

A unique string that identifies the node.

Version	The approximate TIBCO BusinessEvents version, as configured by the connection property. The exact version information is not available via JMX.
Last Update	The date and time the data was last updated in the display.

BE Events

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](#)” 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](#)” display, which shows information for a single BE event.
- **Event Cache Hits:** Opens the “[TIBCO BE Event Cache Hits Table](#)” display, which
- **Event Hit Summary:** Opens the “[TIBCO BE Event Hit Summary](#)” display, which

TIBCO BE Events Table

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](#)” 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	

Filter By:

The display might include these filtering options:

Node: Select a node containing the agent for which you want to view metrics.

Agent ID Select the agent ID for which you want to view metrics.

Fields and Data:

Events The total number of events in the table.

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.

Event The name of the event.

Node The name of the node.

Avg Time Pre RTC s The average amount of time taken for the event to begin its run to completion cycle, in seconds.*

Avg Time in RTC s The average amount of time taken for the event to complete (once it has started) its run to completion cycle, in seconds.*

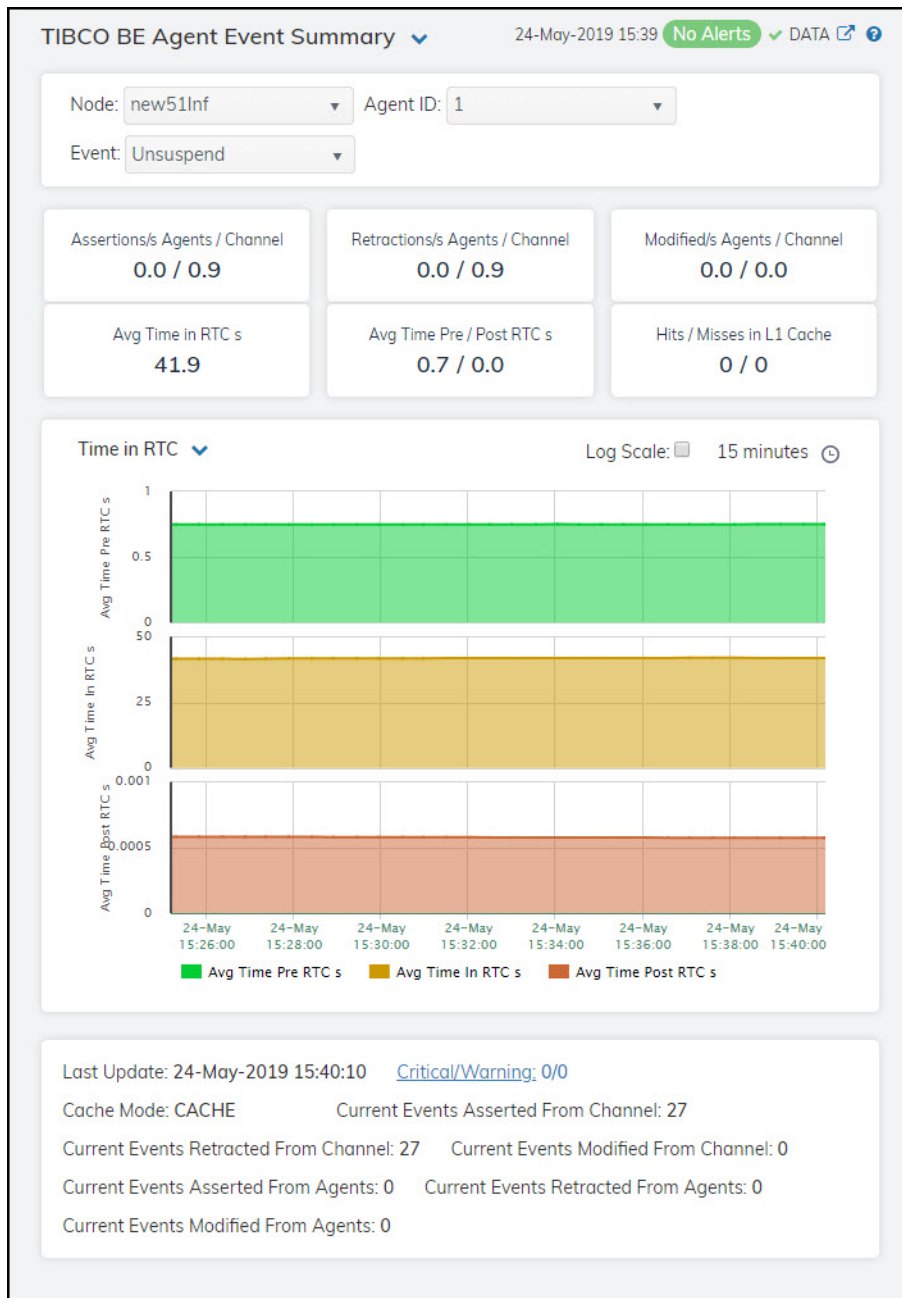
Avg Time Post RTC s	The average amount of time taken by the event after its run to completion cycle has ended., in seconds.*
Cache Mode	Lists the mode used by the event, which can be either CACHE (only) or MEMORY (only).*
Hits in L1 Cache	The number of times data has been searched for in the L1 cache since the last data update.*
Misses n L1 Cache	The number of times data has been searched for in the L1 cache, but was not found, since the last data update.*
Recovered	The number of times data is not found in the L1 cache, but is found in a different cache, since the last data update.*
Asserted from Agents	The number of times the event was asserted by an agent into the Rete network.*
Asserted from Channel	The number of times the event was asserted into the Rete network via the channel.*
Modified from Agents	The number of times the event was modified by an agent in the Rete network.*
Modified from Channel	The number of times the event was modified in the Rete network via the channel.*
Retracted from Agents	The number of times the event was retracted/deleted by an agent from the Rete network.*
Retracted from Channel	The number of times the event was retracted/deleted from the Rete network via the channel.*
L1 Cache Hits/s	The rate of L1 cache hits.
L1 Cache Misses/s	The rate of L1 cache misses.
Recovered / s	The rate of recovered data.
Assertions/ s (Agent)	The rate of event assertions into the Rete network by the agent.
Assertions/ s (Channel)	The rate of event assertions into the Rete network via the channel.
Modified/s (Agent)	The rate of event modifications in the Rete network by the agent.
Modified/s (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 /s (Channel)	The rate of event retractions/deletions from the Rete network via the channel.
agentId	The ID of the agent.

Expired	When checked (true), 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.

TIBCO BE Agent Event Summary

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](#)" 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.



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.

- Node:** Select the node for which you want to view metrics.
- Agent** Select the agent for which you want to view metrics.

Event: Select the event for which you want to view metrics.

Fields and Data:

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.

Assertions/s Agents/Channel	The rate of assertions per second from agent, and the rate of assertions per second from channel.
Retractions/s Agents/Channel	The rate of retractions per second from agent, and the rate of retractions per second from channel.
Modified Agents/Channel	The rate of modified events per second from agent, and the rate of modified events per second from channel.
Avg Time in RTC s	The average time spent in RTC, in seconds.
Avg Time Pre/Post RTC s	The averager time in pre-RTC, and the average time in post-RTC.
Hits/Misses in L1 Cache	The number of hits in the L1 cache, and the number of misses in the L1 cache.

Trend Graphs

Time in RTC

Shows metrics for the selected event:

Avg Time Pre RTC s -- Traces the average amount of time taken for the event to begin its run to completion cycle.

Avg Time in RTC s -- Traces the average amount of time taken for the event to complete (once it has started) its run to completion cycle.

Avg Time Post RTC s -- Traces the average amount of time taken by the event after its run to completion cycle has ended.

Agent Rates

Shows metrics for the selected event:

Assertions/s -- Traces the rate of event assertions, per second, via the agents.

Retractions/s -- Traces the rate of event retractions, per second, via the agents.

Modified/s -- Traces the rate of events modified, per second, via the agents.

Channel Rates

Shows metrics for the selected event:

Assertions/s -- Traces the rate of events asserted, per second, via the channel.

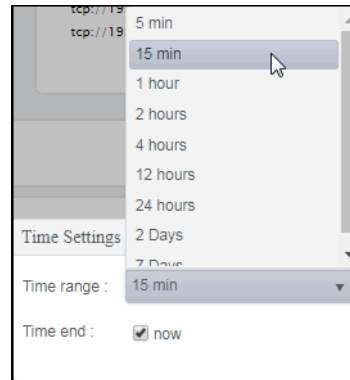
Retractions/s -- Traces the rate of events retracted, per second, via the channel.

Modified/s -- Traces the rate of events modified, per second, via the channel.

Use Rates 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).

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.

Time Settings Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

Last Update	The date and time of the last data update.
Current Events Asserted From Channel	The number of times the event was asserted into the Rete network via the channel since the last data update.*
Current Events Asserted From Agents	The number of times the event was asserted into the Rete network via the agent since the last data update.*
Critical/Warning	The number of critical and warning alerts.
Current Events Retracted From Channel	The number of event retractions/deletions from the Rete network via the channel.
Current Events Retracted From Agents	The number of event retractions/deletions from the Rete network via the agent.
Cache Mode	Lists the mode used by the event, which can be either CACHE (only) or MEMORY (only).*
Current Events Modified From Channel	The number of event modifications in the Rete network via the channel.
Current Events Modified From Agents	The number of event modifications in the Rete network via the agents.

TIBCO BE Event Cache Hits Table

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"](#) 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	

Filter By:

The display might include these filtering options:

Cluster: Select a cluster for which you want to view metrics.

Nodes: Select a node for which you want to view metrics.

Events The total number of events in the table.

Node-Level Event Cache Hits 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.

Event	The name of the event.
Node	The name of the node.
Cache Size	The size of the event's cache.*
Get Avg Time	The average time taken for a "get" event for the node.*
Put Avg Time	The average time taken for a "put" event for the node.*
Remove Avg Time	The average time taken for a "remove" event for the node.*
Gets/s	The rate of "get" operations for the event.
Puts/s	The rate of "put" operations for the event.
Removes/s	The rate of "remove" operations for the event.
Get Count	The total number of "get" operations for the event.*
Put Count	The total number of "put" operations for the event.*
Remove Count	The total number of "remove" operations for the event.*
Num Handles In Store	The number of handles in the Backing Store for the event.*
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.
Time Stamp	The date and time, relative to the Data Server, that data was last collected for the engine.

TIBCO BE Event Hit Summary

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.

**Filter By:**

The display might include these filtering options:

- Cluster:** Select a cluster containing the node and event for which you want to see metrics.
- Nodes:** Select a node containing the event for which you want to see metrics.
- Event** Select the event for which you want to see metrics.

Fields and Data:

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.

Cache Size/Handles in Store	The size of the cache, and the number of handles in the Backing Store for the event.*
Avg Time Gets ms/Puts ms	The average time taken for a "get" operation (in milliseconds), and the average time taken for a "put" operation (in milliseconds).
Gets / Gets/s	The number of "get" operations for the event since the last data update, and the rate of "get" operations for the event.
Puts / Puts/s	The number of "put" operations for the event since the last data update, and the rate of "put" operations for the event.
Removes/Removes/s	The number of "remove" operations for the event since the last data update, and the rate of "remove" operations for the event.
Total Gets/Total Puts	The total number of "get" operations for the event, and the total number of "remove" operations for the event.

Trend Graphs**Hit Rates**

Shows metrics for the selected cluster/node/event combination:

Gets/s -- Traces the rate of "gets" per second for the event.

Puts/s -- Traces the rate of "puts" per second for the event.

Removes/s -- Traces the rate of "removes" per second for the event.

Current Hits

Shows metrics for the selected cluster/node/event combination:

Gets/s -- Traces the number of "gets" for the event.

Puts/s -- Traces the number of "puts" for the event.

Removes/s -- Traces the number of "removes" for the event.


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.

Time Settings Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

Last Update

The date and time in which the data was last updated.

Critical/Warning

The number of critical and warning alerts.

Expired

When **true**, 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.

BE Concepts

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”](#) 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”](#), which shows details and alert status for a BE concept.
- **Channels:** Opens the [“TIBCO BE Channel Status Table”](#) display, which shows information for destinations and channels for a single BE node.
- **Inference Agents Table:** Opens the [“TIBCO BE Inference Agents Table”](#) display, which displays agents data for a specific cluster.
- **RTC Reports Table:** Opens the [“TIBCO BE RTC Txn Manager Reports”](#) display, which displays reports data for a specific cluster.

TIBCO BE Concepts Table

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”](#) 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	

Filter By:

The display might include these filtering options:

- Cluster:** Choose a cluster to see metrics for.
- Nodes:** Choose a node to see metrics for.

Concepts: The total number of concepts in the table.

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.

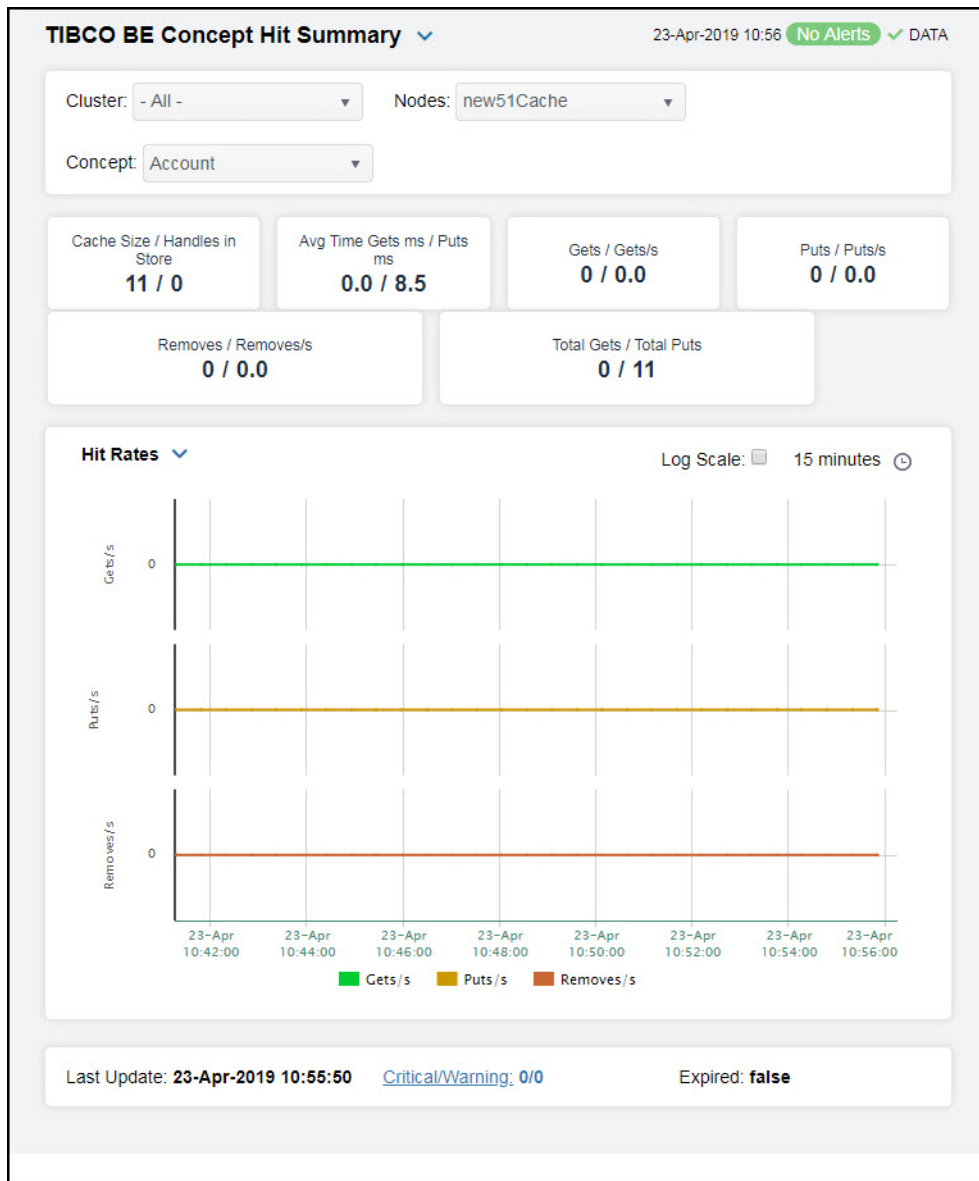
Concept	The name of the concept.
Connection	The name of the connection.
Cache Size	The size of the concept's cache. *
Get Avg Time	The average time taken for a "get" operation. *
Put Avg Time	The average time taken for a "put" operation. *
Remove Avg Time	The average time taken for a "remove" operation. *
Gets/s	The rate of "gets" for the concept.
Puts/s	The rates of "puts" for the concept.
Removes/s	The rate of "removes" for the concept.
Get Count	The total number of "gets" for the concept. *
Put Count	The total number of "puts" for the concept. *
Remove Count	The total number of "removes" for the concept. *
Num Handles In Store	The number of handles in the Backing Store for the concept. *
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.
Time Stamp	The date and time, relative to the Data Server, that data was last collected for the concept.

TIBCO BE Concept Hit Summary

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.



Filter By:

- Cluster:** Select a cluster containing the node and concept for which you want to view metrics.
- Nodes:** Select a node containing the concept for which you want to view metrics.
- Concept** Select the concept for which you want to view 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.

Cache Size/Handles in Store	The size of the cache, and the number of handles in the Backing Store.*
Avg Time Gets ms / Puts ms	The average cache access time taken for a "get" operation (in milliseconds), and the average cache access time taken for a "put" operation (in milliseconds).*
Gets/ Gets/s	The number of cache accesses for "get" operations for the concept, and the rate of "get" operations for the concept.
Puts/ Puts/s	The number of cache accesses for "put" operations for the concept, and the rate of "put" operations for the concept.
Removes/ Removes/s	The number of cache accesses for "remove" operations for the concept, and the rate of "remove" operations for the concept.
Total Gets/ Total Puts	The total number of cache accesses for "get" operations for the concept, and the total number of cache accesses for "put" operations for the concept.

Trend Graphs**Hit Rates**

Shows metrics for the selected cluster/node/concept combination:

Gets/s -- Traces the rate of "get" operations for the concept.

Puts(/sec)-- Traces the rate of "put" operations for the concept.

Removes(/sec)-- Traces the rate of "remove" operations for the concept.

Current Hits

Shows metrics for the selected cluster/node/concept combination:

Gets(/sec) -- Traces the number of "get" operations for the concept.

Puts(/sec)-- Traces the number of "put" operations for the concept.

Removes(/sec)-- Traces the number of "remove" operations for the concept.

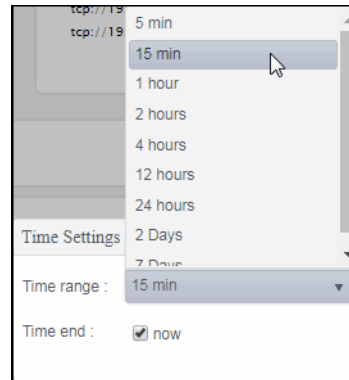
Use Rates

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.

Time Settings Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

Last Update The date and time in which the data was last updated in the display.

Critical/Warning The number of critical and warning alerts.

Expired When true, 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.

TIBCO BE Channel Status Table

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	

Filter By:

The display might include these filtering options:

Cluster: Choose a cluster to see metrics for.

Nodes: Choose a node to see metrics for.

Destinations: The total number of destinations in the table.

Channels: The total number of channels in the table.

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 URI	The Uniform Resource Identifier (URI) for the destination.*
Suspended	Denotes whether the destination is suspended.*
Num Events Received	The number of events received by the destination.*
Number of Events Sent	The number of events sent by the destination.*

Received Events/s	The rate of events received by the destination.
Received Events Last Interval	The rate of events received since the last data update.
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.
Time Stamp	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.

Channel URI	The Uniform Resource Identifier (URI) for the channel.*
State	The current state of 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.

TIBCO BE Inference Agents Table

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"](#) 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

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.

Agents: The number of agents currently in the table.

Inference Agents 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.

Cluster Name	The name of the TIBCO BusinessEvents cluster.
Node	The name of the node.
Agent ID	A unique string that identifies the agent.
Agent Class	The name of the agent's class.
Type	The type of agent (Inference, Cache, Query, or Dashboard).*
Current State	The current state of the agent.*

Started	When checked, denotes that the agent is started.*
Suspended	When checked, denotes that the agent is suspended.*
Concurrent	When checked, denotes that it is a concurrent agent.*
Queue Capacity	The queue capacity for the agent.*
Queue Size	The queue size for the agent.*
Thread Count	The total number of threads for the agent.*
Total # Rules Fired	The total number of rules fired for the agent.*
Rules/s	The rate of rules fired for the agent.
Avg Receive Time	See TIBCO documentation for more information.*
Avg Txn Commit Time	The average amount of time taken to commit a transaction.*
Cache Queue Remaining	The total amount of remaining space on the cache queue.*
DB Ops Queue Remaining	The total amount of remaining space on the DB Operations queue.*
Hit Ratio	See TIBCO documentation for more information.*
Job Rate	See TIBCO documentation for more information.*
L1 Cache Max Size	The maximum size of the L1 cache.*
L1 Cache Size	The current size of the L1 cache.*
Max Active	See TIBCO documentation for more information.*
# Event Threads	The total number of currently active event threads.*
# Jobs	The total number of currently active jobs.*
Priority	See TIBCO documentation for more information.*
Read Only	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 (true), 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.
- Time Stamp** The date and time, relative to the Data Server, that data was last collected for the destination.

TIBCO BE RTC Txn Manager Reports

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”](#) 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	

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.

Reports: The number of reports currently in the table.

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.

Cluster Name	The name of the TIBCO BusinessEvents cluster.
Node	The name of the node.
Avg Action Txn ms	The average amount of time taken for an action transaction, in milliseconds.*
Avg Cache Queue Wait Time ms	The average cache queue wait time, in milliseconds.*
Avg Cache Txn ms	The average amount of time taken for a cache transaction, in milliseconds.*
Avg DB Ops Batch Size	The average database operation batch size.*
Avg DB Queue Wait Time ms	The average database queue wait time, in milliseconds.*
Avg DB Txn ms	The average amount of time taken for a database transaction, in milliseconds.*
Avg Successful Txn Time ms	The average amount of time taken for a successful transaction, in milliseconds.*
Last DB Batch Size	The size of the last database batch.*
Pending Actions	The total number of pending actions.*
Pending Cache Writes	The total number of pending cache writes.*
Pending DB Writes	The total number of pending database writes.*
Pending Events to Ack	The total number of pending events that need to be acknowledged.*
Pending Locks to Release	The total number of pending locks that need to be released.*
Total DB Txns Completed	The total number of database transactions that have been completed.*
Total Successful Txns	The total number of successful transactions.*
Total Errors	The total number of errors.*

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.
Time Stamp	The date and time, relative to the Data Server, that data was last collected for the destination.

TIBCO BusinessWorks

The HTML version features an overview display, ["TIBCO BusinessWorks Overview Display"](#) (pictured below), and the following Views which can be found under **Components** tab > **Middleware**:

- ["BW Applications"](#): The displays in this View present BusinessWorks application performance metrics.
- ["BW Containers"](#): The displays in this View present BusinessWorks container performance metrics.
- ["BW Application Nodes"](#): The displays in this View present BusinessWorks AppNode performance metrics.
- ["BW Application Slices"](#): The displays in this View present BusinessWorks AppSlice performance metrics.
- ["BW Processes"](#): The displays in this View present BusinessWorks process performance metrics.
- ["BW Activities"](#): The displays in this View present BusinessWorks activity performance metrics.

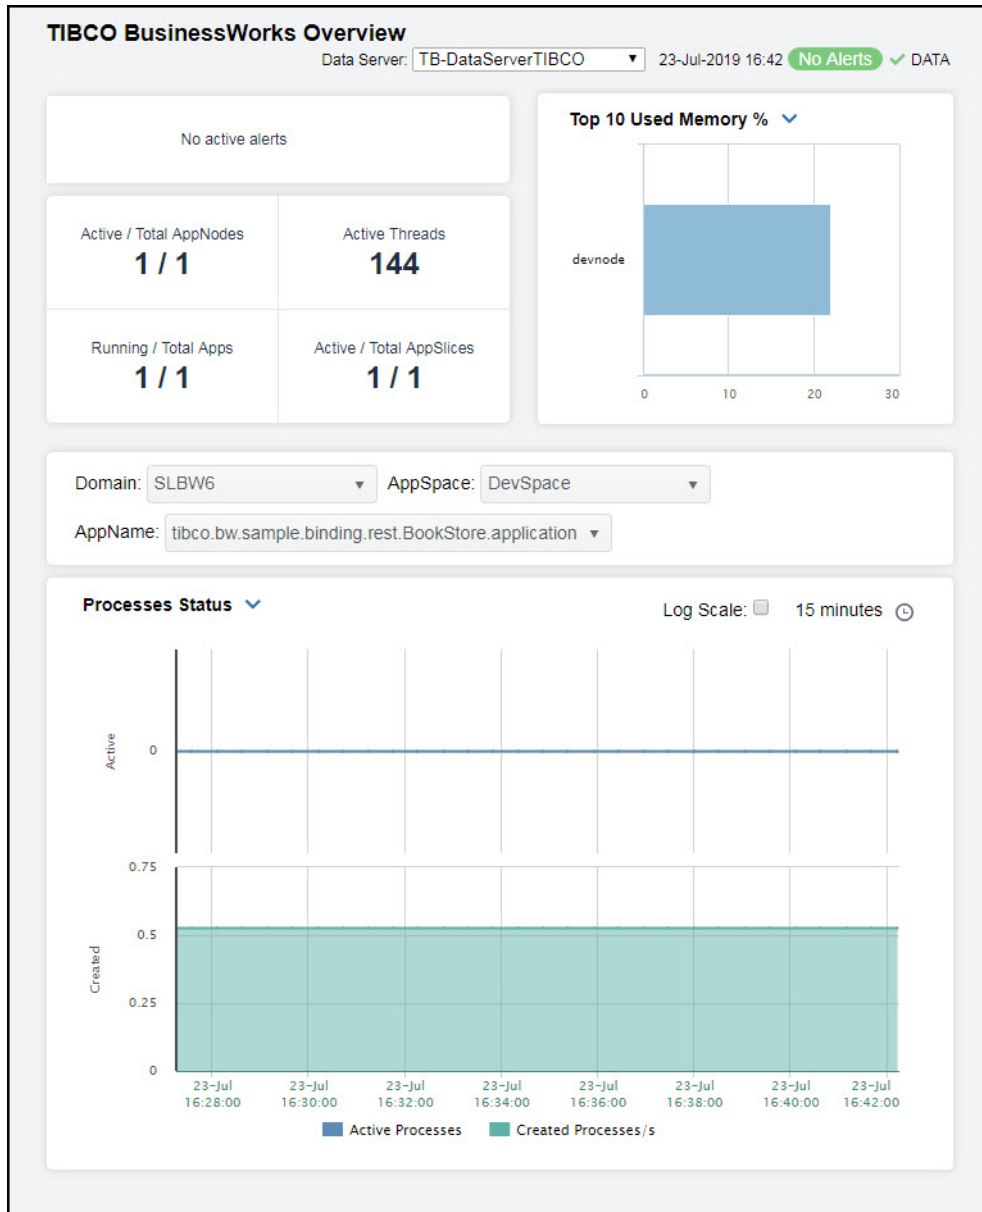
TIBCO BusinessWorks Overview Display

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

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”](#) 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”](#), which shows server and alert status for all BW5 applications.
- **BW Application**: Opens the [“TIBCO BusinessWorks Application Summary”](#) display, which shows information for a single application.

TIBCO BusinessWorks Applications Table

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”](#) 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.

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.applicatio	✓	0

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.

Fields and Data:

- Applications:** The total number of applications in the AppSpace.
- Running** The total number of applications currently running in the AppSpace.

Table:

Each row in the table is a different application.

- Domain** The domain in which the application resides.
- AppSpace** The AppSpace in which the application resides.
- Name** The name of the application.
- 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 application.
State	The current status of the application. Valid values are Running and Stopped .
Deployment	The type of deployment.
AppNodes	The total number of AppNodes associated with the application.
Active Processes	The number of currently active application processes.
Active/s	The rate of processes becoming active per second.
Created Processes	The number of application processes that have been created.
Created/s	The number of application processes created per second.
Completed Processes	The number of completed application processes.
Completed/s	The rate of processes being completed.
Most Recent Exec Time ms	The number of milliseconds for the most recently executed process.
Rate Exec Time ms/s	The number of processes executed, in milliseconds per second.
Suspended Processes	The number of suspended application processes.
Failed Processes	The number of failed application processes.
Version	The application version.
Module	The application module.
Shared Module	The shared module, if any.
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 BusinessWorks > 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 the row data was last updated.

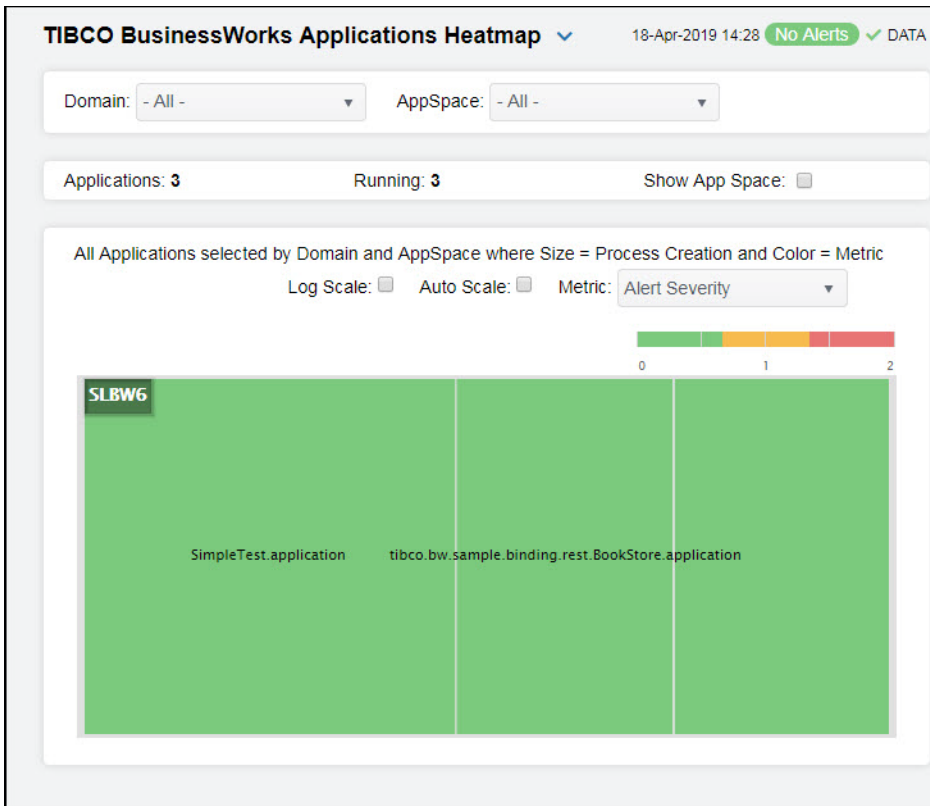
TIBCO BusinessWorks Applications Heatmap

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 AppSpace** 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"](#) display.






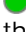







Filter By:





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.

Fields and Data:

- Applications:** The total number of Applications.
- Running** The total number of Applications currently running.
- Show AppSpace** Displays the name of the associated AppSpace in the heatmap when selected.
- 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.

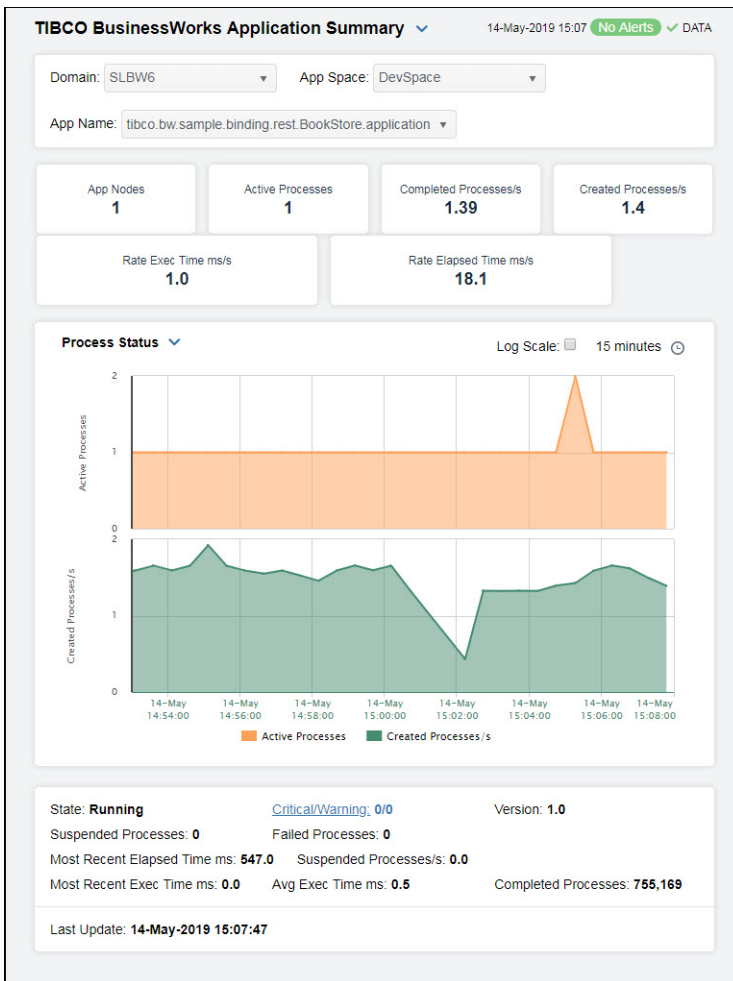
Auto Scale	Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value. NOTE: Some metrics auto-scale automatically, even when Auto is not selected.
Metric	Select the metric driving the heatmap display. The default is Alert Severity . 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" display for a detailed view of metrics for that particular application.
Alert Severity	<p>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:</p> <ul style="list-style-type: none">  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.
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 count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
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 count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.
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 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.
Suspended / sec	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.

Failed / sec	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.
Exec Time / sec	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.
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.
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.

TIBCO BusinessWorks Application Summary

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”](#) display or the [“TIBCO BusinessWorks Processes Table”](#) 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.

**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:

- AppNodes** The number of AppNodes running on this application.
- Active Processes** The number of active processes for this application.
- Completed Processes/s** The rate of completed processes, per second, for this application.
- Created Processes/s** The rate of processes being created, per second, on this application.

Rate Exec Time ms/s	The rate at which the application is accumulating process execution time, in milliseconds per second.
Rate Elapsed Time ms/s	The rate at which the application accumulates process elapsed time, in milliseconds per second.
Trend Graphs	<p>Process Status Traces the sum of process metrics across all processes in all slices of the selected application.</p> <p>Active Processes -- Traces the number of currently active application processes.</p> <p>Created Processes/s -- Traces the rate of created application processes.</p> <p>Process Performance Traces the sum of process metrics across all processes in all slices of the selected application.</p> <p>Rate Exec Time ms/s -- Traces the rate at which the application is accumulating process execution time, in milliseconds per second.</p> <p>Rate Elapsed Time ms/s -- Traces the rate at which the application accumulates process elapsed time, in milliseconds per second.</p>
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	Select a time range from the drop down menu varying from 5 Minutes to Last 7 Days . By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

State	The current status of the application. Valid values are Running and Stopped .
Suspended Processes	The number of suspended application processes.
Suspended Processes/Seconds	The rate of processes being suspended.
Completed Processes	The number of completed processes.
Critical/Warning	The number of critical and warning alerts.
Failed Processes	The number of failed processes.
Most Recent Exec Time ms	The number of milliseconds for the most recently executed process.
Version	The application version.
Most Recent Elapsed Time ms	The most recent elapsed time for a process, in milliseconds.
Avg Exec Time ms	The average number of milliseconds for processes to execute for the selected application.
Last Update	The date and time of the last data update.

BW Containers

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](#)" 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](#)", which is a color-coded heatmap view of selected container performance metrics.
- **BW Container Summary**: Opens the "[TIBCO BusinessWorks Container Summary](#)" display, which shows current and historical metrics for a single container.

TIBCO BusinessWorks 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. Check performance and utilization metrics for processes (for example, **Active Processes**, **Active/s**, **Created Processes**, **Completed Processes**, **Suspended Processes**), execution rates (for example, **Rate Exec Time ms/s**), memory and CPU utilization (for example, **Used Memory** and **Used CPU%**) as well as many others.

By default, all containers are listed in the table, but you can enter a string in the **Application 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"](#) display.

TIBCO BusinessWorks Containers Table

06-Sep-2019 13:23 No Alerts ✓ DATA

Containers: 1 Running: 1 Running Only:

Application Name: * RegEx:




Application Name	Alert Level	Alert Count	State	Contai
tibco.bwce.sample.binding.rest.BookStore.applicat	✓	0	Running	d13a34c

Fields and Data:

- Containers:** 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 table.
- Domain:** Choose a domain to show data for in the display.
- Application Name** Enter a string (all or part of a container name) to filter the data shown in the display. If you enter part of an 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*.
- RegEx** Toggles the **Application 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).

Table:

Each row in the table is a different application.

Application Name	The name of the application.
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 in the container.
State	The current status of the application. Valid values are Running and Stopped .
Container ID	The ID of the container.
Active Processes	The number of currently active processes in the container.
Active/s	The rate of processes becoming active.
Created Processes	The number of processes that have been created since the last data update.
Created/s	The number of processes created per second in the container.
Completed Processes	The number of completed processes in the container.
Completed/s	The rate of processes being completed.
Most Recent Exec Time ms	The number of seconds for the most recently executed process.
Rate Exec Time ms/s	The rate of processes executed in the container, in milliseconds per second.
Suspended Processes	The number of suspended application processes in the container.
Failed Processes	The number of failed processes in the container.
Version	The version of the container.
Module	The name of the container module.
Shared Module	The name of the shared module, if any.
Time Stamp	The date and time the row data was last updated.

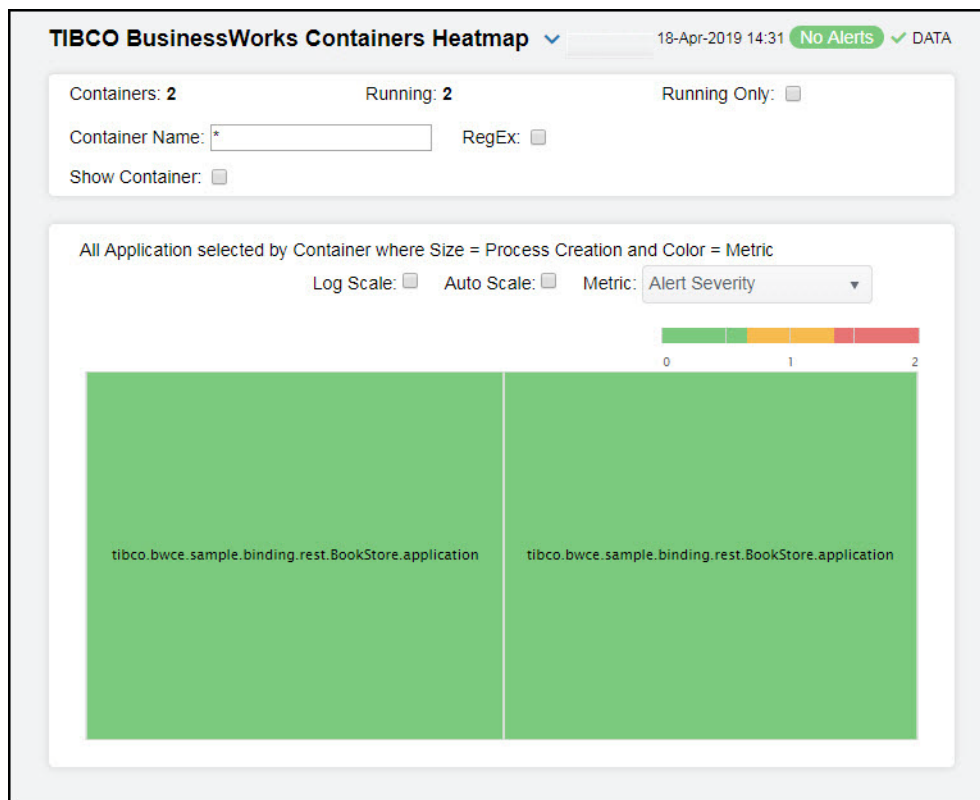
TIBCO BusinessWorks Containers Heatmap

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"](#) display.





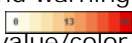











Filter By:



The display might include these filtering options:

Containers:	The total number of containers currently shown in the heatmap.
Running	The total number of containers currently running.
Running Only	Select to show only running containers in the heatmap.
Container Name	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*. You can also enable the RegEx toggle to just enter a portion of the container name.
RegEx	Toggles the Container Name filter to accept Regular Expressions for filtering. For example, if your container name is ContNameOne and this option was toggled on, you could enter "Name" (without using "*" to display the container in the heatmap).
Show Container	Check to include container name labels in the heatmap.

Fields and Data:

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.
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	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" display for a detailed view of metrics for that particular container.
Alert Severity	<p>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:</p> <ul style="list-style-type: none">  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.

Completed	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.
Suspended	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.
Failed	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.
Created/s	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.
Suspended /s	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.
Failed/s	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.
Exec Time/s	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.
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.

Avg 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.
Avg 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.

TIBCO BusinessWorks Container Summary

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"](#) display, where you can view additional container data.

You can toggle between a **Processes** trend graph and a **Resources** trend graph. The **Processes** trend graph shows the number of active containers, the container created rate, the total execution rate, and the total elapsed rate over a selected time range. The **Resources** trend graph shows **CPU%** utilization, current number of **Threads**, **Memory%** utilization and **Bytes** used over a selected time range.

Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display.

**Filter By:**

The display might include these filtering options:

Container Choose the container for which you want to view data in the display.


Fields and Data:

CPU % The percentage of CPU used by the AppNode.

Active Threads The total number of active threads on the AppNode.

Memory % The percentage of memory used by the AppNode.

Total Processes	The total number of processes on the container.
Active Processes	The total number of active processes on the AppNode.
Failed Processes	The total number of failed processes on the AppNode.
Processes Trend Graph	<p>Traces the sum of process metrics across all processes in all slices of the selected container.</p> <p>Active -- Traces the number of currently active application processes on the container.</p> <p>Created/s -- Traces the number of created application processes on the container.</p> <p>Total Execution ms/s -- Traces the rate at which the application is accumulating process execution time, in milliseconds per second, on the container.</p> <p>Total Elapsed ms/s -- Traces the rate at which the application accumulates process elapsed time, in milliseconds per second, on the container.</p>
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	Select a time range from the drop down menu varying from 5 Minutes to Last 7 Days . By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

Name	The name of the application.
Up Since	The length of time that the application has been up and running.
Completed Processes	The number of completed processes.
Critical/Warning	The number of critical and warning alerts.
Created Processes	The number of created processes.
Completed/s	The rate of completed processes.
State	The current status of the application. Valid values are Running and Stopped .
Created/s	The rate at which processes are being created.
Free Memory MB	The amount of free memory, in megabytes.
Last Update	The date and time of the last data update.

BW Application Nodes

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”](#) 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”](#), which shows a color-coded heatmap view of utilization metrics.
- **BW Application Node**: Opens the [“TIBCO BusinessWorks Application Node Summary”](#) display, which shows Current and historical metrics for a single AppNode.

TIBCO BusinessWorks Application Nodes Table

View BusinessWorks data shown in the [“TIBCO BusinessWorks Application Nodes 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. Drill-down and investigate by clicking a row to view details for the selected AppNode in the [“TIBCO BusinessWorks Application Node Summary”](#) display.

TIBCO BusinessWorks Application Nodes Table 17-Apr-2019 14:29 No Alerts DATA

Domain: - All - AppSpace: - All -

AppNodes: 2

Domain	AppSpace	AppNode	Alert Level	Alert Count	State	Host
SLBW6	DevSpace	testnode	✔	0	ACTIVE	slhost9agent(SLBW6)
SLBW6	DevSpace	devnode	✔	0	ACTIVE	qawin5agent(SLBW6)

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.

AppNodes: The total number of rows in the table.

Table:

Column values describe the AppNode.

Domain The domain in which the AppNode resides.

AppSpace The AppSpace in which the AppNode resides.

AppNode The name of the AppNode.

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 AppNode.

State The current status of the application. Valid values are **Running** and **Stopped**.


Host	The host on which the AppNode resides.
Deployment	The type of deployment.
Processes	The number of processes running.
Active Threads	The number of currently active threads.
Total Memory	The total amount of used and free memory, in megabytes.
Used Memory	The amount of used memory, in megabytes.
Free Memory	The amount of free memory, in megabytes.
Used Memory %	The percent (%) used memory.
Used CPU %	The percent (%) used CPU.
Active Processes	The number of currently active application processes.
Active/s	The rate of application processes becoming active.
Created Processes	The number of application processes that have been created.
Created /s	The number of application processes created per second.
Completed Processes	The number of completed application processes.
Completed/s	The rate of application processes being completed.
Most Recent Exec Time ms	The number of seconds for the most recently executed process, in milliseconds.
Rate Exec Time ms/s	The number of application processes executed per second.
Suspended Processes	The number of suspended application processes.
Failed Processes	The number of failed application processes.
System Process ID	A unique string identifier for the process.
Up Since	The date and time the AppNode was last started.
Time Stamp	The date and time the row data was last updated.
Source	Name of RTView Data Server sending this data (or localhost).
Expired	When checked (true), 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 BusinessWorks > 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.

TIBCO BusinessWorks Application Nodes Heatmap

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.

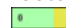






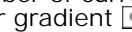

Use the **Show AppSpace** check-box to include or exclude labels in the heatmap. Mouse over a rectangle to see additional metrics. By default, this display shows **Alert Severity**. Drill-down and investigate by clicking a rectangle in the heatmap to view details for the selected application in the ["TIBCO BusinessWorks Application Node Summary"](#) display. Toggle between the commonly accessed **Table** and **Heatmap** displays by clicking the drop down list on the display title.





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.
- AppNodes:** The total number of AppNodes in the AppSpace.

Show AppSpace	When selected, the names of the AppSpaces associated with the AppNodes display in the heatmap.
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.
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	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 " display for a detailed view of metrics for that particular AppNode.
Alert Severity	<p>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:</p> <ul style="list-style-type: none">  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 maximum CPU used percentage in the heatmap. The middle value in the gradient bar indicates the average amount.
Memory Used%	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 0 to the maximum memory used percentage in the heatmap. The middle value in the gradient bar indicates the average amount.
Active Processes	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 0 to the maximum number of active processes 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 number of created processes 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 creation rate in the heatmap. The middle value in the gradient bar indicates the average count.
Failed Processes	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 0 to the maximum number of failed processes in the heatmap. The middle value in the gradient bar indicates the average number.

TIBCO BusinessWorks Application Node Summary

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](#)" 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.

TIBCO BusinessWorks Application Node Summary

18-Apr-2019 14:37 No Alerts DATA

Domain: standalone AppSpace: 1dd3b7b21c65
AppNode: 1dd3b7b21c65

CPU % 9.8	Active Threads 117	Memory % 17.0	Total Processes 3
Active Processes 0	Failed Processes 0		

Process totals by AppNode and App where Color = Metric
Metric: Alert Severity

System Utilization Log Scale: 5 minutes

State: ACTIVE	Critical/Warning: 0/0	Up Since: 0d 22:42
Created Processes: 2,181	Completed Processes: 2,181	Free Memory MB: 2,438,361,088
Last Update: 18-Apr-2019 14:37:14		

Filter By:

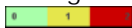




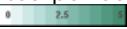

The display might include these filtering options:

- 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:

- CPU %** The percent (%) CPU used on the AppNode.
- Active Threads** The number of currently active threads for the AppNode.
- Memory %** The percentage of memory utilization on the AppNode.
- Total Processes** The total number of processes on the AppNode.
- Active Processes** The total number of active processes on the AppNode.
- Failed Processes** The total number of failed processes on the AppNode.

Heatmap Metrics Each rectangle in the heatmap represents an AppNode. Click a rectangle to drill-down and investigate in the [“TIBCO BusinessWorks Application Summary”](#) 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.
- 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.
- Created/s** 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 creation rate in the heatmap. The middle value in the gradient bar indicates the average count.
- Average Execution** The average number of processes executed 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 number of processes executed in the heatmap. The middle value in the gradient bar indicates the average count.

Trend Graphs	<p>System Utilization Traces the sum of process metrics across all processes for all applications on the AppNode. CPU% -- Traces the percent (%) CPU used on the AppNode. Threads -- Traces the number of threads.</p> <p>Memory Utilization Traces the sum of process metrics across all processes for all applications on the AppNode. Memory Percentage -- Traces the percentage of memory used. Bytes -- Traces the total number of bytes.</p>
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	Select a time range from the drop down menu varying from 5 Minutes to Last 7 Days . By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

State	The current status of the application. Valid values are Running and Stopped .
Created Processes	The total number of processes that have been created on the AppNode.
Critical/Warning	The number of critical and warning alerts.
Completed Processes	The total number of processes that have been completed on the AppNode.
Up Since	The number of days, hours, minutes, and seconds that the AppNode has been up and running.

Free Memory MB	The amount of available memory on the AppNode, in megabytes.
Last Update	The date and time of the last data update.

BW Application Slices

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"](#) 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"](#), which shows process execution metrics for all AppSlices.
- **BW Application Slice**: Opens the ["TIBCO BusinessWorks Application Slice Summary"](#) display, which shows current and historical metrics for a single AppSlice.

TIBCO BusinessWorks Application Slices Table

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"](#) 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

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.




Fields and Data:

- AppSlices:** The total number of AppSpaces listed in the table.
- Running** The total number of applications currently running in the AppSpace.

Table:

Each row in the table is a different AppNode.

- Domain** The domain in which the AppSpace resides.
- AppSpace** The AppSpace the AppNode is associated with.
- AppNode** The name of the selected AppNode.
- Name** The name of the application.

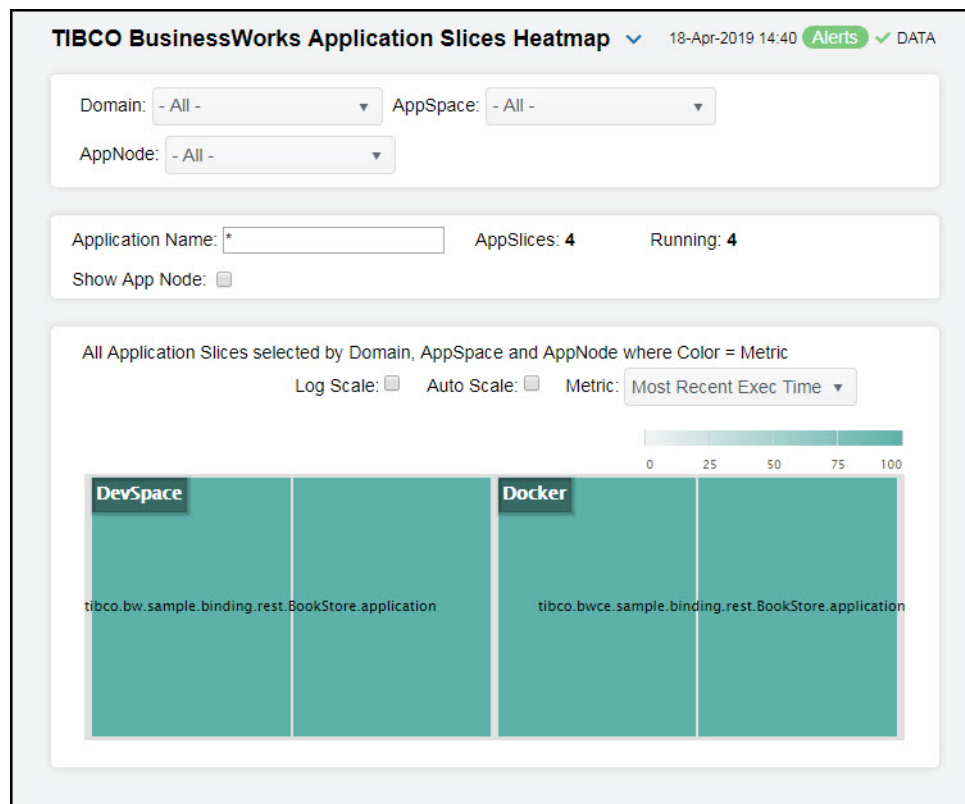
Alert Level	The most critical alert state for alerts in the row:  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	The total number of active alerts for the AppNode.
State	The current status of the application. Valid values are Running and Stopped .
Active Processes	The number of currently active application processes.
Active/s	The rate of application processes becoming active, per second.
Created Processes	The number of application processes that have been created.
Created/s	The number of application processes created per second.
Completed Processes	The number of completed application processes.
Completed/s	The number of application processes completed per second.
Most Recent Exec Time ms	The number of seconds for the most recently executed process.
Rate Exec Time ms/s	The number of processes executed per second.
Suspended Processes	The number of suspended application processes.
Failed Processes	The number of failed application processes.
Expired	When checked (true), 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 BusinessWorks > 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 the row data was last updated.

TIBCO BusinessWorks Application Slices Heatmap

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"](#) 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 to limit data shown in the display.

AppNode: Choose an AppNode to show data for in the display.

Fields and Data:

AppSlices The number of AppSlices in the display.

Running The total number of AppSlices currently running in the display.


Running Only Select to show only running applications in the display.


Show App Node Check to include labels in the heatmap.


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.


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 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"](#) 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.

Created/s	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.
Suspended /s	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.
Failed/s	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.
Exec Time/s	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.
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 execution time in an AppSlice in the heatmap. The middle value in the gradient bar indicates the average execution time.
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 average execution time in the heatmap. The middle value in the gradient bar indicates the average execution time.
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 average elapsed time in the heatmap. The middle value in the gradient bar indicates the average elapsed time.

TIBCO BusinessWorks Application Slice Summary

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”](#) 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](#)" 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: AppSpace:

AppNode:

Application Name:

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:

Process Status Log Scale: 15 minutes

State: Running	Critical/Warning: 0/0	Version: 1.0
Suspended Processes: 0	Failed Processes: 0	Average Elapsed Time ms: 14
Suspended Processes/s: 0.0	Failed Processes/s: 0.0	Avg Exec Time ms: 0
Completed: 755,727		

Last Update: **14-May-2019 15:13:23**

Filter By:

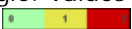




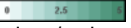

The display might include these filtering options:

Domain:	Select a domain for which you want to view data in the display.
AppSpace	Select an AppSpace for which you want to view data in the display.
AppNode:	Select an AppNode for which you want to view data in the display.
Application Name:	Select an Application Name for which you want to view data in the display.

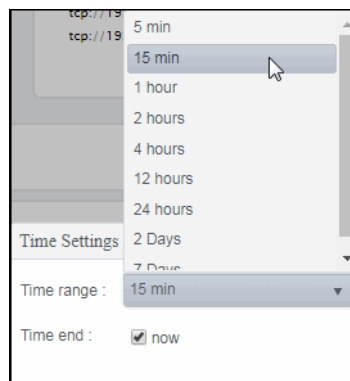
Fields and Data:

Processes	The total number of processes on the AppSlice.
Active Processes	The total number of active processes on the AppSlice.
Completed Processes/s	The rate of completed processes, per second, on the AppSlice.
Created Processes/s	The rate of created processes, per second, on the AppSlice.
Rate Exec Time ms/s	The time spent on execution of processes, in milliseconds per second, in the AppSlice.
Rate Elapsed Time ms/s	The rate of elapsed time spent for processes, in milliseconds per second.

Heatmap Metrics Each rectangle in the heatmap represents an AppSlice. Click a rectangle to drill-down and investigate in the [“TIBCO BusinessWorks Process Summary”](#) display.

Alert Severity	<p>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:</p> <ul style="list-style-type: none">  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	<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 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.</p>
Created/s	<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 0 to the maximum creation rate in the heatmap. The middle value in the gradient bar indicates the average count.</p>
Average Execution	<p>The average number of processes executed 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 number of processes executed in the heatmap. The middle value in the gradient bar indicates the average count.</p>

Trend Graphs	<p>Process Status Traces the sum of process metrics across all processes for all applications on the AppNode.</p> <p>Active Processes -- Traces the number of active processes.</p> <p>Created/s -- Traces the number of processes created per second.</p> <p>Process Performance Traces the sum of process metrics across all processes for all applications on the AppNode.</p> <p>Exec Time/s -- Traces the rate at which the application accumulates process execution time, in milliseconds per second</p> <p>Elapsed Time/s -- Traces the rate at which the application is accumulating process elapsed time, in milliseconds per second.</p>
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	Select a time range from the drop down menu varying from 5 Minutes to Last 7 Days . By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

State	The current status of the application. Valid values are Running and Stopped .
Suspended Processes	The total number of suspended processes.
Suspended Processes/s	The rate of processes being suspended.
Completed Processes	The total number of completed processes summed across all processes in one AppSlice of the application.

Critical/Warning	The number of critical and warning alerts.
Failed Processes	The total number of failed processes.
Failed Processes/s	The number of failed application processes per second.
Version	The application version.
Average Elapsed Time ms	The average amount of elapsed time for processes, in milliseconds.
Average Exec Time ms	The average number of milliseconds for processes to execute.
Last Update:	The date and time the data was last updated.

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. Clicking **BW Processes** from the left/navigation menu opens the "[TIBCO BusinessWorks Processes Table](#)" 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](#)", which shows process execution metrics for all processes.
- **BW Process:** Opens the "[TIBCO BusinessWorks Process Summary](#)" display, which shows current and historical metrics for a single process.

TIBCO BusinessWorks Processes Table

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](#)” 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

Filter By:

The display might include these filtering options:

- Domain:** Select a domain for which you want to view data in the display.
- AppSpace** Select an AppSpace for which you want to view data in the display.
- AppNode:** Select an AppNode for which you want to view data in the display.
- Application** Select an Application for which you want to view data in the display.




Fields and Data:

- Processes:** The total number of processes in the AppSpace, which are listed in the table.

Table:

Each row in the table is a different process. Column values are associated with the process.

- Domain** The domain in which the process resides.

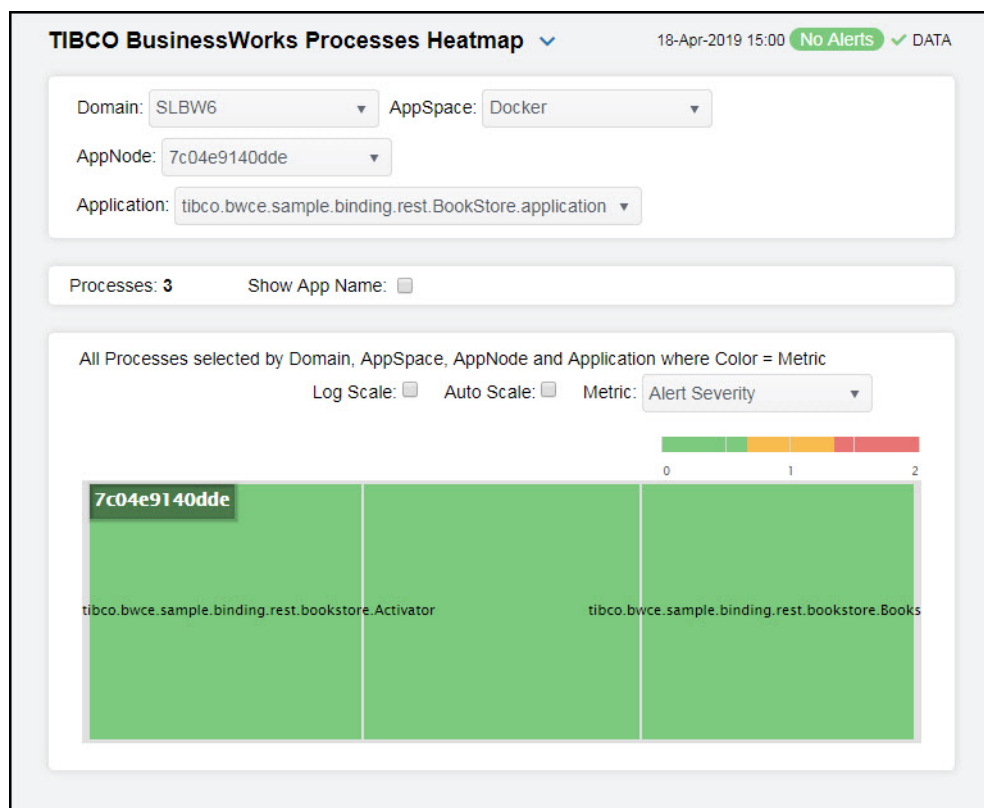
AppSpace	The AppSpace in which the process resides.
AppNode	The AppSpace in which the process resides.
Application Name	The name of the application in which the process is running.
Process Name	The name of the process.
Alert Level	The most critical alert state for alerts in the row:  Red indicates that 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 ms	Total execution time (in milliseconds) for all successfully completed process instances.
Current Total Exec Time	Execution time accumulated during the current polling period.
Rate Exec Time ms/s	Rate execution time in milliseconds per second.
Most Recent Exec Time ms	Execution time (in milliseconds) of the most recently completed process instance.
Total Elapsed Time	Total elapsed time (in milliseconds) for all successfully completed process instances.
Current Total Elapsed Time	Elapsed time accumulated during the current polling period.
Rate Elapsed Time ms/s	Delta elapsed time, in milliseconds per second.
Most Recent Elapsed Time ms	Elapsed clock time (in milliseconds) of the most recently completed process instance.
Active	The number of currently active processes
Created	The number of processes created.
Suspended	The number of process suspensions.
Failed	The number of process failures.
Completed	The number of completed processes.
Current Active	The number of active processes since the last data update.
Active/s	The number of active processes per second.
Current Created	The number of created processes since the last data update.
Created/s	The number of created processes per second.
Current Suspended	The number of suspended processes since the last data update.

Suspended/s	The number of suspended processes per second.
Current Completed	The number of completed processes since the last data update.
Completed/s	The number of completed processes per second.
Current Failed	The number of failed processes since the last data update.
Failed/s	The number of failed processes per second.
Min Exec Time ms	Execution time (in milliseconds) of the process instance that has completed in the shortest amount of execution time.
Max Exec Time ms	Execution time (in milliseconds) of the process instance that has completed in the longest amount of execution time.
Average Exec Time ms	Average execution time (in milliseconds) for all successfully completed process instances.
Min Elapsed Time ms	Elapsed clock time (in milliseconds) of the process instance that has completed in the shortest amount of elapsed time.
Max Elapsed Time ms	Elapsed clock time (in milliseconds) of the process instance that has completed in the longest amount of elapsed time.
Average Elapsed Time ms	Average elapsed clock time (in milliseconds) for all successfully completed process instances.
Count Since Reset	The number of times the process has executed since statistics were reset.
Main Process	The name of the main process.
Version	The version of the application.
Module Name	The name of the application module.
Module Version	The version of the module.
Source	Name of RTView Data Server sending this data (or localhost).
Time Stamp	The date and time the row data was last updated.

TIBCO BusinessWorks Processes Heatmap

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"](#) display.

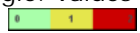















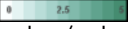


Filter By:

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** Choose an AppName to show data for in the display.

Fields and Data:

Processes:	The total number of processes currently shown in the display.
Show App Name	Check to display the application names in the heatmap.
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.
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 Scale is not selected.
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 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 " display for a detailed view of metrics for that particular process.
Alert Severity	<p>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:</p> <ul style="list-style-type: none">  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	<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 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.</p>
Active Count	<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 0 to the maximum count of active processes in the heatmap. The middle value in the gradient bar indicates the average count.</p>
Completed Count	<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 0 to the maximum count of completed processes in the heatmap. The middle value in the gradient bar indicates the average count.</p>
Suspended Count	<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 0 to the maximum count of suspended processes in the heatmap. The middle value in the gradient bar indicates the average count.</p>
Failed Count	<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 0 to the maximum count of failed processes in the heatmap. The middle value in the gradient bar indicates the average count.</p>

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 number of processes created in the heatmap. The middle value in the gradient bar indicates the average number of created processes.
Suspended / sec	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 number of suspended processes per second in the heatmap. The middle value in the gradient bar indicates the average rate of suspended processes.
Failed / sec	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 rate of failed processes in the heatmap. The middle value in the gradient bar indicates the average count.
Exec Time / sec	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 the heatmap. The middle value in the gradient bar indicates the average execution time rate.
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 most recent execution time in the heatmap. The middle value in the gradient bar indicates the average time.
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 average execution time in the heatmap. The middle value in the gradient bar indicates the average time.
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 most recent elapsed time in the heatmap. The middle value in the gradient bar indicates the average most recent elapsed time.
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 average elapsed time in the heatmap. The middle value in the gradient bar indicates the average elapsed time.

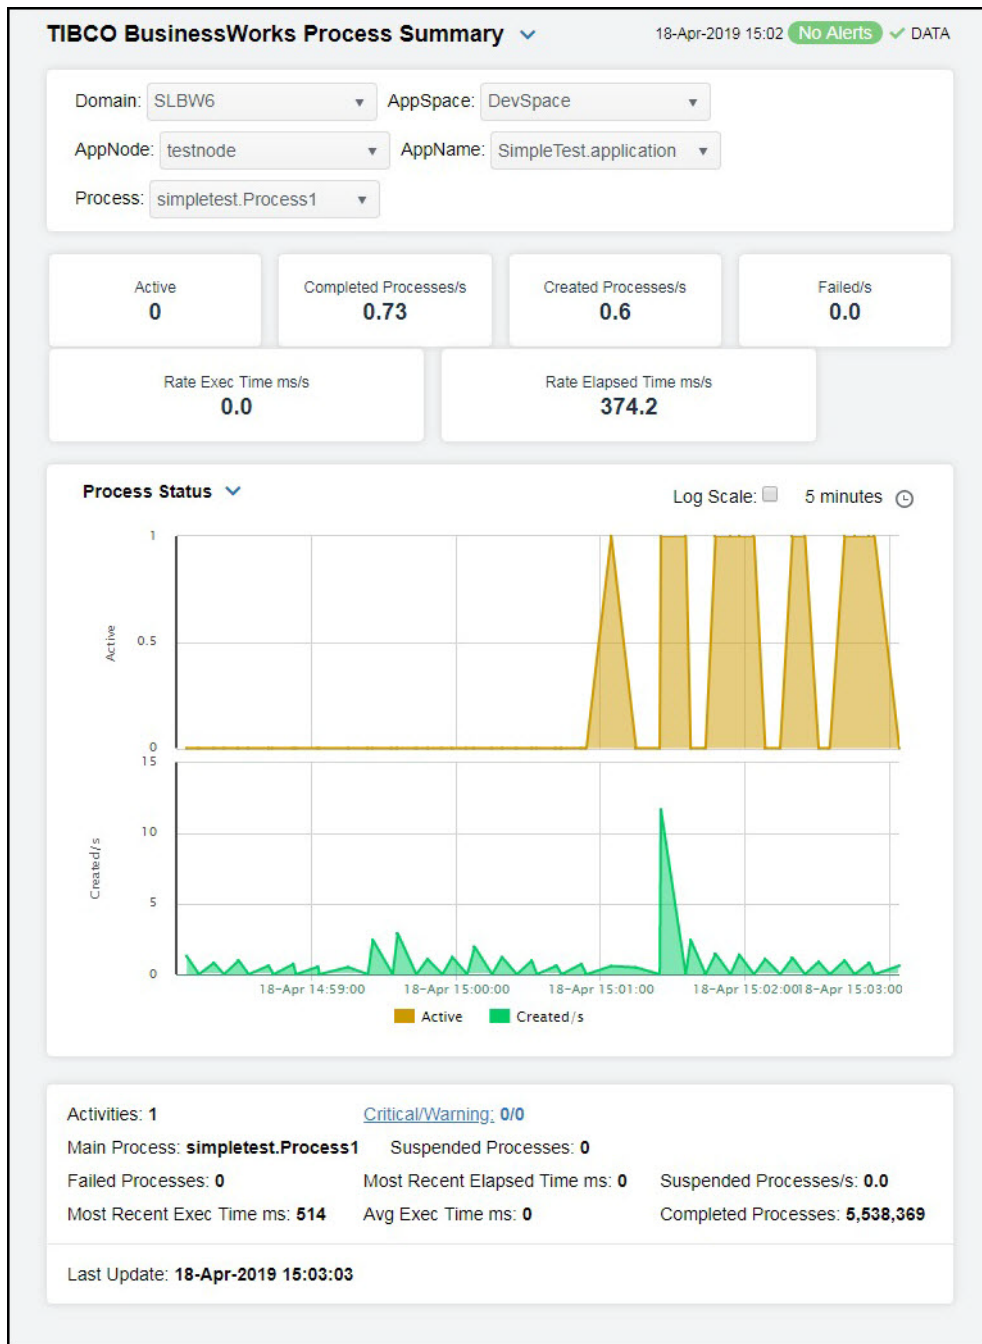
TIBCO BusinessWorks Process Summary

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](#)" 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.



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.
- AppNode:** Select the AppNode for which you want to view data in the display.

AppName Select the application for which you want to view data in the display.

Process Select the process for which you want to view data in the display.

Fields and Data:

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.

Completed Processes/s The number of process instances completed per second.

Created Processes/s The number of process instances created per second.

Failed/s The number of errors accumulated per second.

Rate Exec Time ms/s The amount of execution time accumulated in milliseconds per second.

Rate Elapsed Time ms/s The amount of elapsed time accumulated in milliseconds per second.

Trend Graphs

Process Status

Traces application process and activity metrics for the selected process.

Active -- Traces the number of currently active processes.

Created/s -- Traces the rate of created processes, per second.

Process Performance

Traces application process and activity metrics for the selected process.

Elapsed Time/ -- Traces the rate at which the application is accumulating process elapsed time, in milliseconds per second.


Exec Time/s -- Traces the rate at which the application is accumulating process execution 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.

Time Settings

Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

Activities	The number of activities defined for the process.
Main Process	The name of the main process.
Failed Processes	The number of failed process instances.
Most Recent Exec Time ms	The most recent execution time of any process instance, in milliseconds.
Critical/Warning	The number of critical and warning alerts.
Suspended Processes	The number of suspended processes.
Most Recent Elapsed Time ms	The most recent elapsed time of any process instance, in milliseconds.
Avg Exec Time ms	The average execution time for all completed process instances, in milliseconds.
Suspended Processes/s	The rate of processes being suspended, per second.
Completed Processes	The number of completed processes.
Last Update	The date and time of the last data update.

BW Activities

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"](#) 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"](#), which shows process execution metrics for all activities.
- **BW Activity**: Opens the ["TIBCO BusinessWorks Activity Summary"](#) display, which shows current and historical metrics for a single activity.

TIBCO BusinessWorks Activities Table

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"](#) 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

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.
- AppNode:** Select the AppNode for which you want to view data in the display.
- AppName** Select the application for which you want to view data in the display.
- Process** Select the process for which you want to view data in the display.
- Activity Name Filter** Enter a string to limit data shown in the display.
- Activities** The number of activities displayed in the heatmap as a result of the defined filters.

Table:

Each row in the table is a different process. 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.
Activity Name	The name of the activity.
Application Version	The version of the application.
Execution Count	The total number of process executions.
Error Count	The total number of errors accumulated.
Last Return Code	Status code (OK, DEAD, or ERROR) returned by most recent execution of this activity.
Execution Time	The amount of time used to execute processes.
Elapsed Time	The amount of elapsed time when executing processes.
Min Elapsed Time	The shortest amount of elapsed time.
Max Elapsed Time	The longest amount of elapsed time.
Min Execution Time	The shortest amount of time needed to execute a process.
Max Execution Time	The longest amount of time needed to execute a process.
Most Recent Exec Time ms	The time spent in execution since the last data update.
Most Recent Elapsed Time ms	The elapsed time accumulated since the last data update.
Exec Count Since Reset	Total number of process executions since the last data update.
Current Executions	The number of process executions since the last data update.
Execs/s	The rate of executions, per second, for the activity.
Current Elapsed Time	The elapsed time accumulated since the last data update.
Elapsed Time/s	The elapsed time spent, in milliseconds per second, for the activity.
Current Exec Time	The time spent in execution since the last data update.
Exec Time/s	The rate of executions, per second, for the activity.
Current Errors	

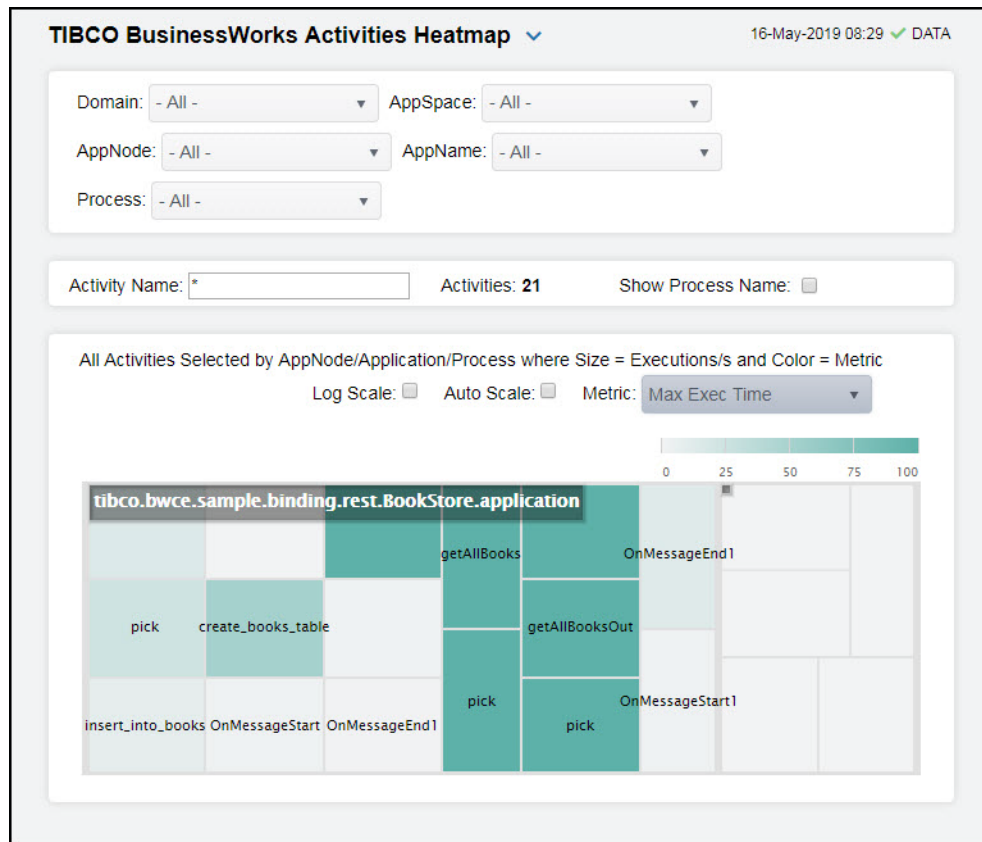
Errors/s	Number of errors accumulated since the last data update.
Avg Exec Time	The average execution time, in milliseconds, for the activity.
Avg Elapsed Time	The average elapsed time, in milliseconds, for the activity.
Avg Errors	The average number of errors for the activity.
Time Stamp	The date and time of the last data update.
Source	Name of RTView Data Server sending this data (or localhost).

TIBCO BusinessWorks Activities Heatmap

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](#)" 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.

**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.
- AppNode:** Select the AppNode for which you want to view data in the display.
- AppName** Select the application for which you want to view data in the display.
- Process** Select the process for which you want to view data in the display.
- Activity Name Filter** Enter a string to limit data shown in the display.
- Activities** The number of activities displayed in the heatmap as a result of the defined filters.
- Show Process Names** Select this check box to display the names of the processes in their respective rectangles in the heatmap.
- 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.

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	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 " display for a detailed view of metrics for that particular activity.
Alert Severity	<p>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:</p> <ul style="list-style-type: none">  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.
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 errors in the heatmap. The middle value in the gradient bar indicates the average error count.
Errors/sec	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.
Exec Count	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.
Exec Time/sec	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.
Most Recent Exec Time	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.
Max Exec Time	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

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"](#) 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
1.1

Errors/s
0.0

Rate Exec Time ms/s
0.5

Avg Exec Time ms
0.5

Rate Elapsed Time ms/s
38.8

Avg Elapsed Time ms
36.3

Utilization Log Scale: 15 minutes 🕒

Execution Count: **42,818**

Error Count: **0**

Min Execution Time: **0**

Min Elapsed Time: **0**

Exec Count Since Reset: **0**

Avg Errors: **0.0**

Current Exec Time: **5**

Current Elapsed Time: **363**

Current Executions: **10**

Current Errors: **0**

Max Execution Time: **2,011**

Max Elapsed Time: **17,317**

Last Update: **16-May-2019 08:32:04**

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.
- AppNode:** Select the AppNode for which you want to view data in the display.
- AppName** Select the application for which you want to view data in the display.
- Process** Select the process for which you want to view data in the display.
- Activity Name** Select the activity name for which you want to view data in the display.

Fields and Data

- Executions /s** The rate of executions, per second, for the activity.
- Errors/s** The rate of errors being generated, per second, for the activity.
- Rate Exec Time ms/s** Time spent in execution, in milliseconds per second, for the activity.
- Avg Exec Time ms** The average execution time, in milliseconds, for the activity.
- Rate Elapsed Time ms/s** The elapsed time spent, in milliseconds per second, for the activity.
- Avg Elapsed Time ms** The average elapsed time, in milliseconds, for the activity.

Trend Graphs

Utilization

Executions -- Traces the number of executions for the activity.

Average Errors -- Traces the average number of errors for the activity.

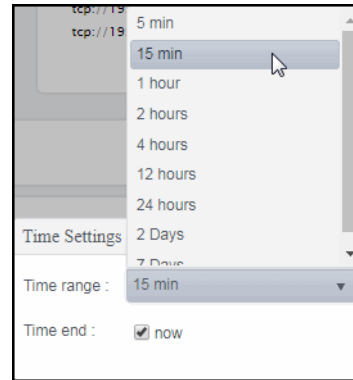
Performance

Average Elapsed Time ms -- Traces the average elapsed time, in milliseconds, for the activity.

Average Exec Time ms -- Traces the average execution time, in milliseconds, for the activity.

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 Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

Execution Count	The total number of process executions.
Error Count	The total number of errors accumulated.
Exec Count Since Reset	Total number of process executions since the last data update.
Avg Errors	The average number of errors for the activity.
Current Executions	The number of process executions since the last data update.
Current Errors	Number of errors accumulated since the last data update.
Min Execution Time	The shortest amount of time needed to execute a process.
Min Elapsed Time	The shortest amount of elapsed time.
Current Exec Time	The time spent in execution since the last data update.

Current Elapsed Time	The elapsed time accumulated since the last data update.
Max Execution Time	The longest amount of time needed to execute a process.
Max Elapsed Time	The longest amount of elapsed time.
Last Update	The date and time of the last data update.

TIBCO BusinessWorks 5 Monitor

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”**: The displays in this View present BusinessWorks 5.0 server performance metrics.
- **“BW5 Engines”**: The displays in this View present BusinessWorks 5.0 engine performance metrics.
- **“BW5 Processes”**: The displays in this View present BusinessWorks 5.0 process performance metrics.
- **“BW5 Activities”**: The displays in this View present BusinessWorks 5.0 activity performance metrics.

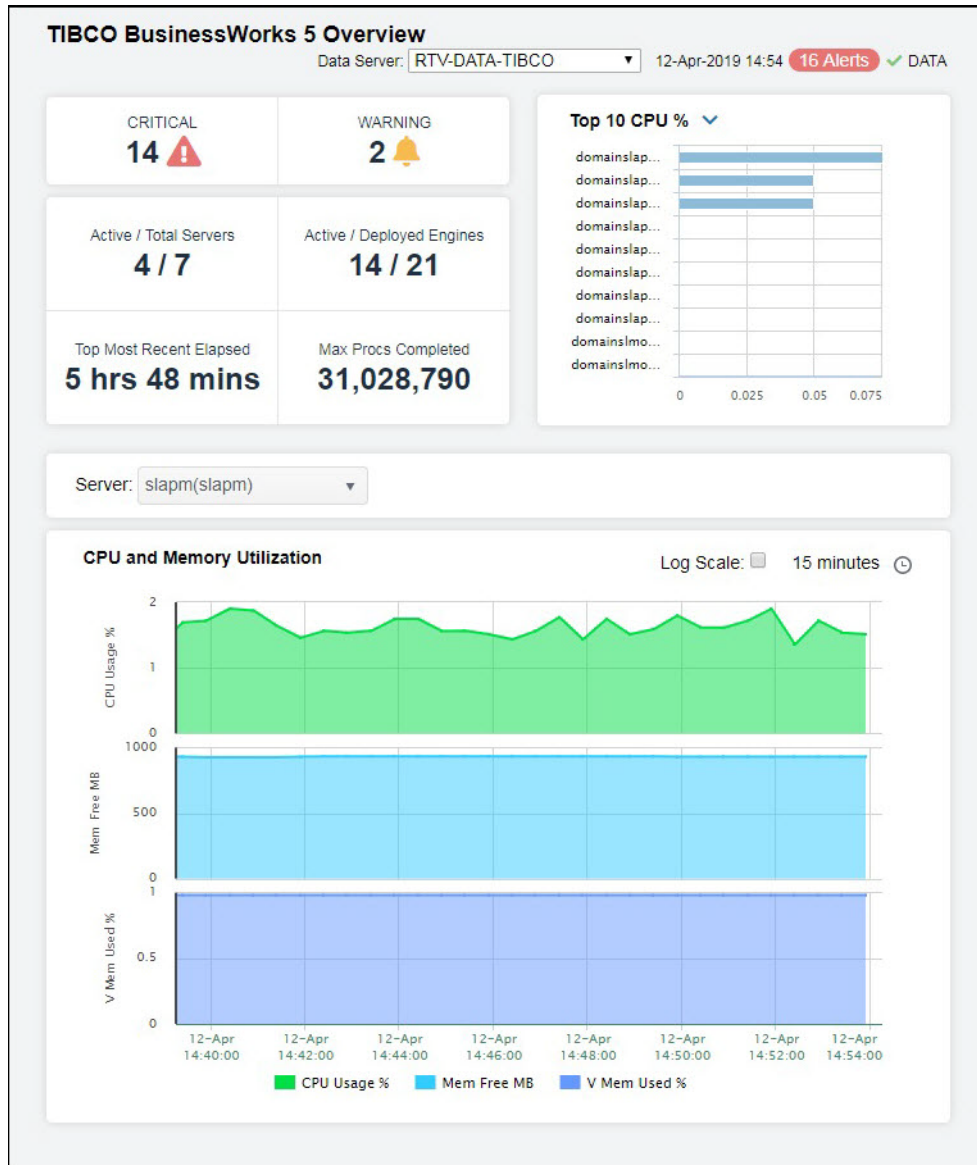
TIBCO BusinessWorks 5 Overview Display

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

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](#)” 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](#)”, which shows server and alert status for all BW5 servers.
- **BW5 Server Summary**: Opens the “[TIBCO BusinessWorks 5 Server Summary](#)” display, which shows information for a single BW5 server.

TIBCO BusinessWorks 5 Servers Table




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](#)” 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.

Server	Alert Level	Alert Count	CPU Usage %	Mem Free MB	V Mem Used %	BW
SLHOST5(domainslhost5)	✓		15.2	1,882.3	0.61	v5.
QAWIN5(-)	✓		55.2	198.9	50.56	
SLHOST21(dev)	✓		9.2	5,456.4	28.52	

Fields and Data

Servers: The number of servers listed in the table.

Table:

Server	Name of Server Agent.
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 application.
CPU Usage %	Percent of server CPU in use.
Mem Free MB	Available physical memory (MB) remaining.
V Memory Used %	Percent of virtual memory used.
BW Version	The TIBCO BusinessWorks version currently in use on the server.
Source	Name of RTView Data Server sending this data (or localhost).
Deployed Engines	Total number of engines deployed on the server.
Active Engines	Number of engines currently active.
Expired	When checked, data has not been received from this host in the specified amount of time.
Time Stamp	Time this data was retrieved.

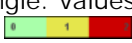
TIBCO BusinessWorks 5 Servers Heatmap







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](#)" 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.



Fields and Data:

- Servers:** The total number of servers 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.
- 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** 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"](#) display for a detailed view of metrics for that particular server.
- 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 BwServersCpuUsedHigh , which is 100. The middle value in the gradient bar indicates the middle value of the range (the default is 50).
V(irtual) Memory Used%	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 0 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 50).
Free Memory	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 0 to the amount of free memory in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
Deployed Engines	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 0 to the number of deployed engines in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
Active Engines	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 0 to the number of active engines in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

TIBCO BusinessWorks 5 Server Summary

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"](#) display, the ["TIBCO BusinessWorks 5 Engines Table"](#) display, or the ["TIBCO BusinessWorks 5 Processes Table"](#) display (depending on which box you select), 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.



Filter By:

The display might include these filtering options:

Server: Select the server for which you want to see data.

Fields and Data

CPU Usage %	Percent of server CPU in use. Click to drill-down to view additional details in the " TIBCO BusinessWorks 5 Servers Table " display.
Memory Free MB	Available physical memory remaining (in MB). Click to drill-down to view additional details in the " TIBCO BusinessWorks 5 Servers Table " display.
Virtual Memory Usage %	Percent of virtual memory used. Click to drill-down to view additional details in the " TIBCO BusinessWorks 5 Servers Table " display.
Deployed Engines	Number of engines currently active. Click to drill-down to view additional details in the " TIBCO BusinessWorks 5 Engines Table " display.
Active Engines	Shows data for the server. Click to drill-down to view additional details in the " TIBCO BusinessWorks 5 Engines Table " display.
Total Running Processes	The number of running processes on this server across all engines. Click to drill-down to view additional details in the " TIBCO BusinessWorks 5 Processes Table " display.

CPU and Memory Utilization Trend Graph

CPU Usage % -- Traces the percentage of server CPU in use.

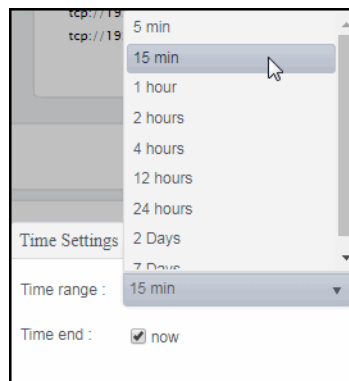
Mem Free MB -- Traces the available physical memory remaining (in MB).

V Mem Used % -- Traces the percentage of virtual memory 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.

Time Settings Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

BW Version	The currently deployed version of TIBCO BusinessWorks 5.
Status	Server status: ACTIVE or EXPIRED.
Critical/Warning	The number of critical and warning alerts.
Threads	The number of threads.
Last Update	The date and time of the last data update.

BW5 Engines

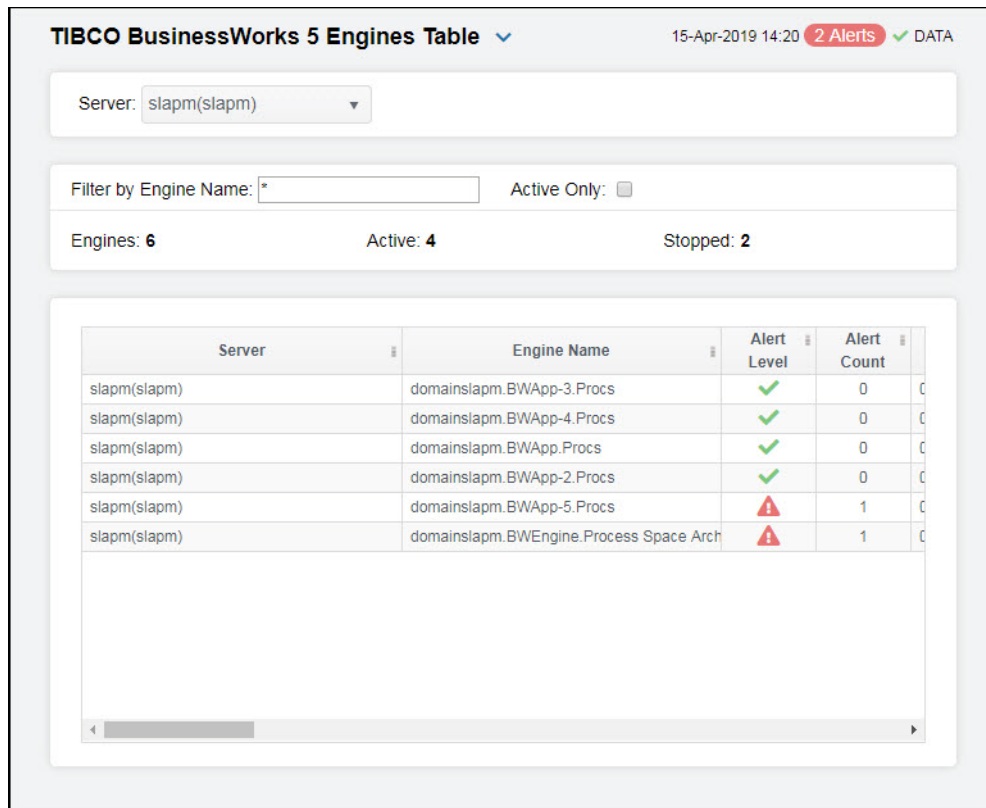
These displays present performance metrics and alert status for all BW5 engines. Clicking **BW5 Engines** from the left/navigation menu opens the "[TIBCO BusinessWorks 5 Engines Table](#)" display, which shows all available utilization metrics for all BW5 engines. The options available under **BW5 Engines** are:

- **All Engines Heatmap**: Opens the "[TIBCO BusinessWorks 5 Engines Heatmap](#)", which shows engine and alert status for all BW5 servers.
- **BW5 Engine Summary**: Opens the "[TIBCO BusinessWorks 5 Engine Summary](#)" display, which shows information for a single BW5 engine.

TIBCO BusinessWorks 5 Engines Table

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"](#) 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: slapm(slapm)

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

Filter By:




The display might include these filtering options:

Server: Select the server for which you want to show data in the display.

Filter by Engine Name: Enter all or part of engine name to view specific engines.

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.
Engines	Number of engines currently being displayed.
Active	Number of engines currently active.
Stopped	The number of stopped engines.

Table:

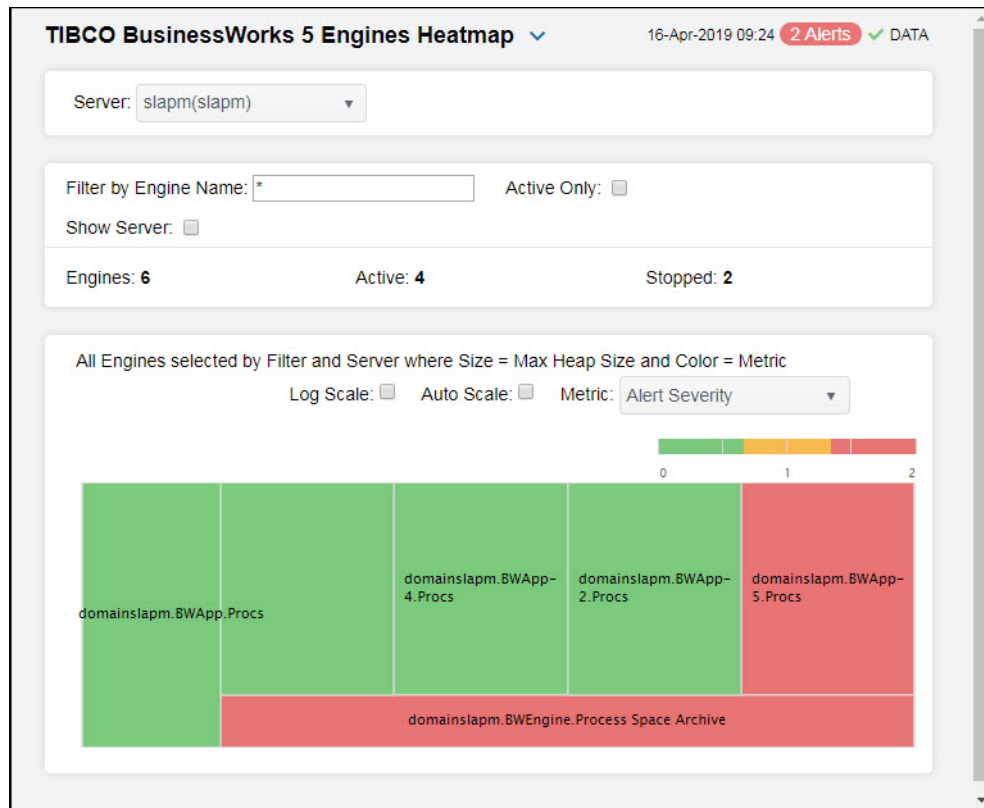
Server	The name of the Server agent.
Engine Name	The name of the engine.
Alert Level	The most critical alert state for alerts in the row:  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	Number of current alerts
Micro Agent Instance	Unique ID of the microagent reporting the metrics.
Process ID	Process ID of engine as recognized by the server.
State	Engine status: ACTIVE, STOPPED, LIMITED, etc. (See All Servers Grid).
Uptime	Uptime in milliseconds since the engine was started.
Threads	Number of threads used by the engine.
CPU %	Percent of server CPU used by engine.
Running Processes	Number of running processes.
Total Bytes	Maximum heap memory this JVM has used.
Free Bytes	Amount of available memory from within the JVM.
Used Bytes	Total bytes of memory within the JVM currently used by the engine. Equal to value of: MaxBytes minus FreeBytes.
Mem Usage KB	Server memory, in KB, used by the engine.
Mem Used %	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.
Total Errors	Total number of errors since the engine was started.
Current Errors	Current number of new errors.
Errors/s	The rate off errors occurring.
Created/s	The rate of processes being created.
Completed/s	The rate of processes being completed.
Active/s	The rate of processes becoming active.

Aborted/s	The rate of processes being aborted.
Deployment	The name of the deployment.
Domain	The name of the domain.
Max Heap Size Bytes	Maximum heap allocated to this engine for the JVM, in bytes.
BW Version	Engine project version number.
Source	Name of RTView Data Server sending this data (or localhost).
Host	Host name of server.
Time Stamp	The date and time of the last data update.

TIBCO BusinessWorks 5 Engines Heatmap











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"](#) 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.

**Filter By:**

The display might include these filtering options:

- Server:** Select the server for which you want to show data in the display.
- Filter by Engine Name:** Enter all or part of engine name to view specific engines.
- 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.
- Show Server** Select this check box to display the associated server names in the heatmap.
- Engines** Number of engines currently being displayed.
- Active** Number of engines currently active.
- Stopped** The number of stopped engines.
- 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.

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	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 " display for a detailed view of metrics for that particular engine.
Alert Severity	<p>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:</p> <ul style="list-style-type: none">  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).
Memory Used%	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 0 to the alert threshold of BwEngineMemUsedHigh , which is 100 . The middle value in the gradient bar indicates the middle value of the range (the default is 50).
Active Processes	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 0 to the maximum count of active processes in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
Running Processes	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 0 to the maximum count of running processes in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
Created Processes	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 0 to the maximum count of created processes in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

- 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 rate of processes created per second in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
- 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 errors in the heatmap. The middle value in the gradient bar indicates the average alert count.

TIBCO BusinessWorks 5 Engine Summary

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”](#) display or to the [“TIBCO BusinessWorks 5 Processes Table”](#) 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 memory utilization section contains two line charts. The top chart shows 'Current Mem MB' on the y-axis (0 to 300) against time on the x-axis (16-Apr 10:30:00 to 16-Apr 10:34:00). The bottom chart shows 'Used Mem MB' on the y-axis (0 to 150) against the same time period. A tooltip is displayed over the data at 10:32:15, indicating a Current Mem MB of 226.74 and a Used Mem MB of 44.45.

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**

Filter By:

The display might include these filtering options:

Server: Select the server for which you want to view data in the display.

Engine: Select the engine for which you want to view data in the display

Fields and Data

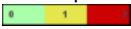







CPU %	The percentage of CPU utilization on the engine.
Total Threads	The total number of threads being executed on the engine.
Memory Used %	The percentage of memory utilization on the engine.
Total Processes	The total number of deployed processes on the engine.
Running Processes	The total number of running processes on the engine.
Total Errors	The total number of errors generated on the engine.

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.

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 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 Process Summary"](#) display for a detailed view of metrics for that particular engine.

Alert Severity	<p>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:</p> <ul style="list-style-type: none">  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	<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 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.</p>
Created Processes	<p>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 0 to the maximum count of created processes in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>
Completed Processes	<p>The 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 created processes in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>
Avg Execution	<p>The average execution time for 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 created processes in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>

**Trend
Graphs****System Utilization**

CPU % -- Traces the CPU utilization for the engine.

Running Processes -- Traces the number of running processes on the engine.

Memory Utilization

Current Mem MB -- Traces the current memory available for the engine.

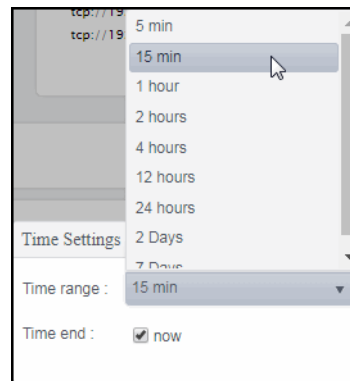
Used Mem MB -- Traces the used memory on the engine.

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**

Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

State

ACTIVE -- The BW microagent is providing live data and the engine is assumed active.

SUSPENDED -- This state is reported by the BW microagent.

STANDBY -- This state is reported by the BW microagent.

STOPPING -- This state is reported by the BW microagent.

STOPPED -- This state is reported by the BW microagent.

LIMITED -- Live data has been received from TIBCO, but deployment data from the custom RTView MicroAgent has not been received.

EXPIRED -- The associated server for the engine is currently in an EXPIRED state and is unavailable or stopped sending data.

**Free Memory
MB**

The amount of free memory, in megabytes.

Total Created Processes	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.
Total Aborted Processes	The total number of aborted processes.
Critical/Warning	The number of critical and warning alerts.
Domain	The name of the domain.
Total Completed Processes	The total number of completed processes.
Uptime	Days, hours, and minutes since the engine was started.
BW Version	
Last Update	The date and time of the last data update.

BW5 Processes

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](#)" 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](#)", which shows process execution metrics for all BW Engines.
- **BW5 Process Summary**: Opens the "[TIBCO BusinessWorks 5 Process Summary](#)" display, which shows historical and current metrics for a single process, including average execution times and execution counts.

TIBCO BusinessWorks 5 Processes Table

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](#)" 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	Com
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	

Filter By:

The display might include these filtering options:

Server: Select the server for which you want to view data in the display.

Engine: Select the engine for which you want to view data in the display.

Processes The number of processes listed in the table.

Table:

Process 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 Number of current alerts

Total CPU Total CPU usage in percent.

Created/s Change in Created per second.

Completed/s Change in Completed per second.

Created Number of process instances created for this process definition.

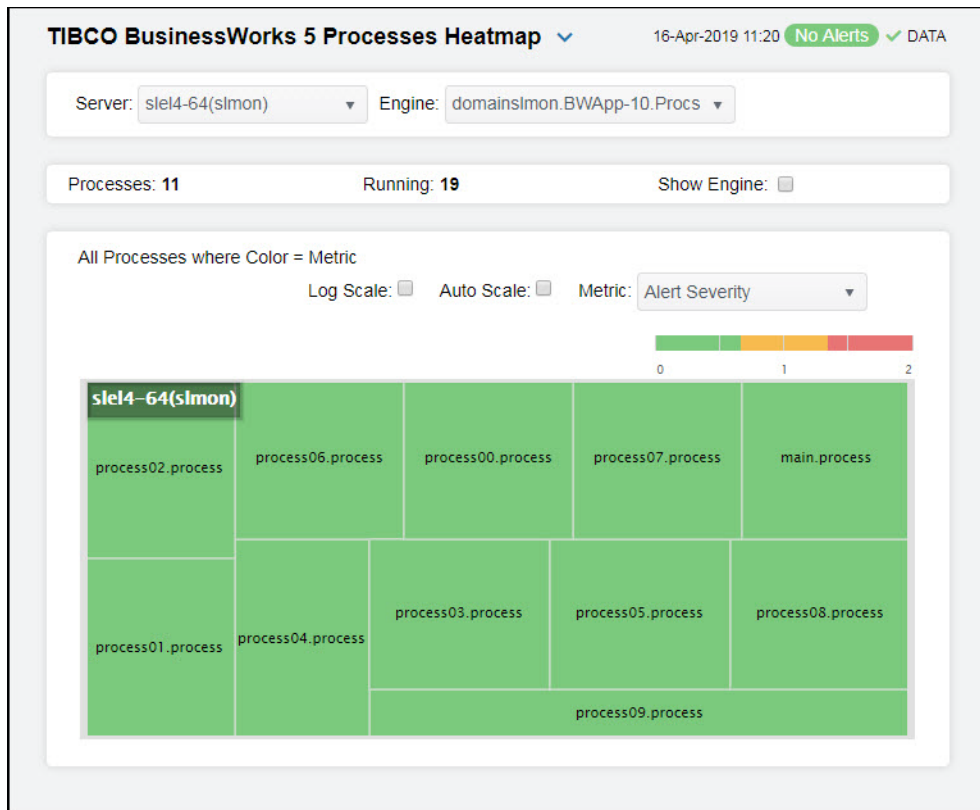
Completed	Number of process instances successfully completed.
Total Exec Time ms	Total execution time (in milliseconds) for all successfully completed process instances.
Exec Time/s	Delta execution time per second.
Current Total Execs	Execution time accumulated during the current polling period.
Min Exec Time ms	Execution time (in milliseconds) of the process instance that has completed in the shortest amount of execution time.
Max Exec Time ms	Execution time (in milliseconds) of the process instance that has completed in the longest amount of execution time.
Avg Exec Time ms	Average execution time (in milliseconds) for all successfully completed process instances.
Most Recent Exec Time ms	Execution time (in milliseconds) of the most recently completed process instance.
Total Elapsed Time ms	Total elapsed time (in milliseconds) for all successfully completed process instances.
Current Elapsed ms	Elapsed time accumulated during the current polling period.
Elapsed Time/s	Delta elapsed time per second.
Min Elapsed Time ms	Elapsed clock time (in milliseconds) of the process instance that has completed in the shortest amount of elapsed time.
Max Elapsed Time ms	Elapsed clock time (in milliseconds) of the process instance that has completed in the longest amount of elapsed time.
Avg Elapsed Time ms	Average elapsed clock time (in milliseconds) for all successfully completed process instances.
Most Recent Elapsed Time	Elapsed clock time (in milliseconds) of the most recently completed process instance.
Aborted	Number of times process instances have been aborted.
Current Aborted	Change in Aborted this update.
Aborted/s	Change in Aborted per second.
Queued	Number of times process instances have been queued for execution.
Current Queued	Change in Queued this update.
Queued/s	Change in Queued per second.
Suspended	Number of times process instances have been suspended.
Current Suspended	Change in Suspended this update.
Suspended/s	Change in Suspended per second.
Checkpointed	Number of times process instances have executed a checkpoint.
Current Checkpointed	Change in Checkpointed this update.
Checkpointed/s	Change in Checkpointed per second.

Swapped	Number of times process instances have been swapped to disk.
Current Swapped	Change in Swapped this update.
Swapped/s	Change in Swapped per second.
Time Since Last Update	Time since the last update.
Domain	Name of TIBCO Domain.
Starter	Name of the process starter for the process.
Micro Agent Instance	Unique ID of the microagent reporting the metrics.
Count Since Reset	Number of process instances that have completed since the last reset of the statistics.
Source	Name of RTView Data Server sending this data (or localhost).
Time Stamp	Time of last update.

TIBCO BusinessWorks 5 Processes Heatmap












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"](#) 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.

**Filter By:**

The display might include these filtering options:

- Server:** Select the server for which you want to view data in the display.
- Engine:** Select the engine for which you want to view data in the display.
- Processes** The total number of processes in the display.
- Running** Number of processes currently running.
- Show Engine** Select this check box to display the names of the engines above their respective rectangles in the heatmap.
- 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.
- 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 Scale** is not selected.

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 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" display for a detailed view of metrics for that particular process.
Alert Severity	<p>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:</p> <ul style="list-style-type: none">  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 BwProcessTotalCpuPercentHigh , which is 100 . The middle value in the gradient bar indicates the middle value of the range (the default is 50).
Completed	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 completed processes in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
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 active processes in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
Aborted	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 0 to the maximum count of aborted processes in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
Suspended	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 suspended processes in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
Exec Time/s	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 0 to the maximum execution rate of processes in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

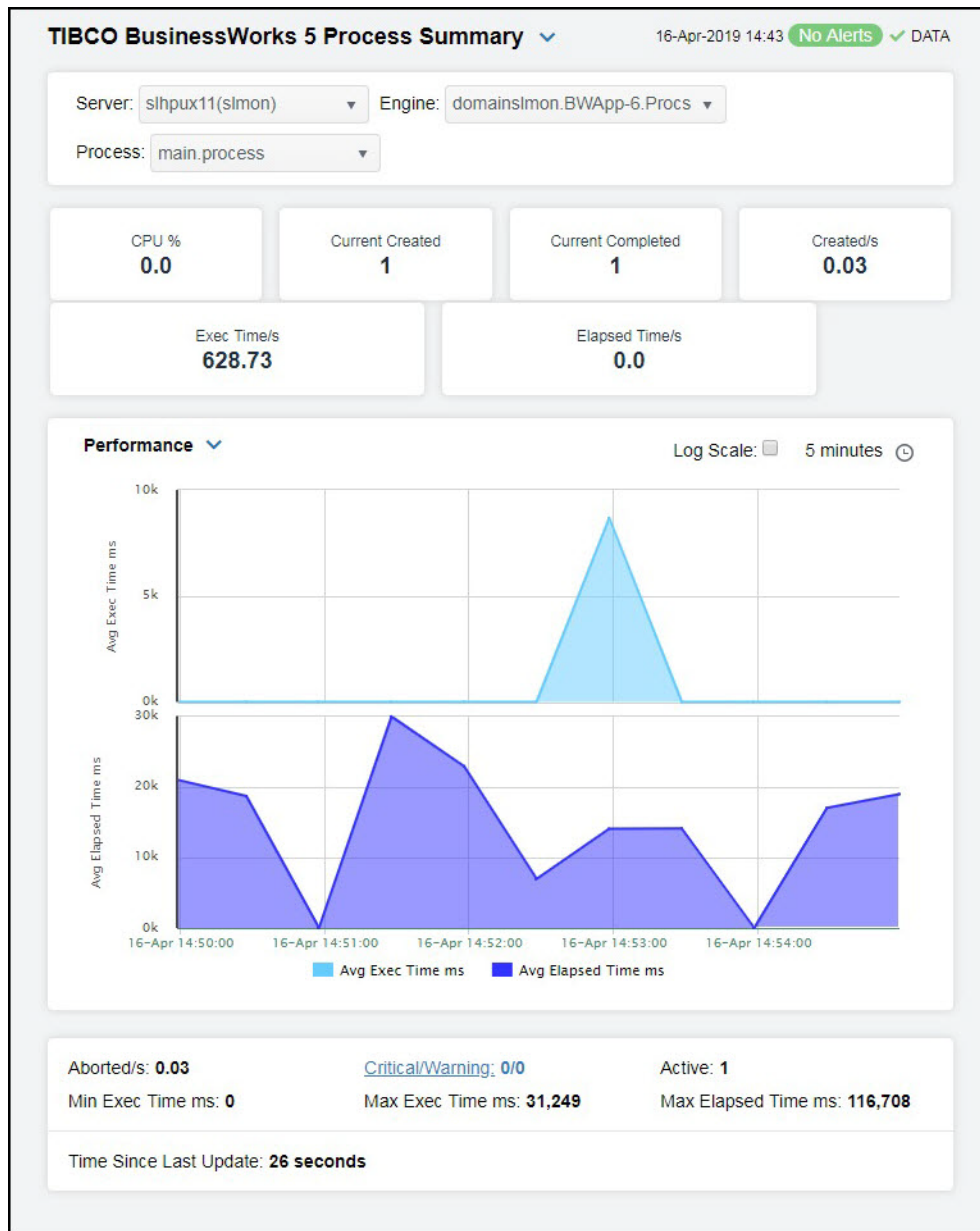
Created/s	The number of processes created per second in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum created rate of processes in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
Aborted/s	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 aborted rate of processes in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
Suspended /s	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 suspended rate of processes in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
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 most recent execution time of processes in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
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 average execution time of processes in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
Most Recent Elapsed Time	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 most recent elapsed time of processes in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
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 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

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"](#) 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.

**Filter By:**

The display might include these filtering options:

Server: Select the server for which you want to view data.


Engine: Select the engine for which you want to view data.

Fields and Data

CPU % The percentage of CPU used by the process.

Current Created The number of processes created since the last data update.

Current Completed	The number of completed processes since the last data update.
Created/s	The rate of processes created, per second.
Exec Time/s	The rate of process execution time, per second.
Elapsed Time/s	The rate of elapsed time, per second, for the process.
Trend Graphs	<p>Utilization</p> <p>CPU % -- Traces the CPU utilization percentage.</p> <p>Created Procs/s -- Traces the number of created processes per second.</p> <p>Performance</p> <p>Avg Exec Time ms -- Traces the average time taken to execute processes, in milliseconds.</p> <p>Avg Elapsed Time ms -- Traces the average elapsed time for processes, in milliseconds.</p> <p>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.</p> <p>Time Settings Select a time range from the drop down menu varying from 5 Minutes to Last 7 Days. By default, the time range end point is the current time.</p>



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

Abprted/s The rate of aborted processes.

Min Exec Time ms	Shortest execution time of any process instance, in milliseconds.
Critical/Warning	The number of critical and warning alerts.
Max Exec Time ms	Longest execution time of any process instance, in milliseconds.
Active	The number of active processes.
Max Elapsed Time ms	Longest elapsed time of any process instance, in milliseconds.
Time Since Last Update	Time since the last update to file of statistics.

BW5 Activities

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”](#) 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”](#), which shows process execution metrics for all activities.
- **BW5 Activity Summary**: Opens the [“TIBCO BusinessWorks 5 Activity Summary”](#) display, which shows historical and current performance metrics for a single activity, including average execution times and execution counts.

TIBCO BusinessWorks 5 Activities Table

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”](#) 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

Filter By:

The display might include these filtering options:

- Server:** Select the server for which you want to view data in the display, or select **All** to view data for all servers.
- Engine:** Select the engine for which you want to view data in the display, or select **All** to view data for all engines.
- Process:** Select the process for which you want to view data in the display, or select **All** to view data for all processes.

Fields and Data

- OK** Number of processes with a status code of OK.
- Error** Number of processes with a status code of Error.
- Dead** Number of processes with a status code of Dead.

Table:

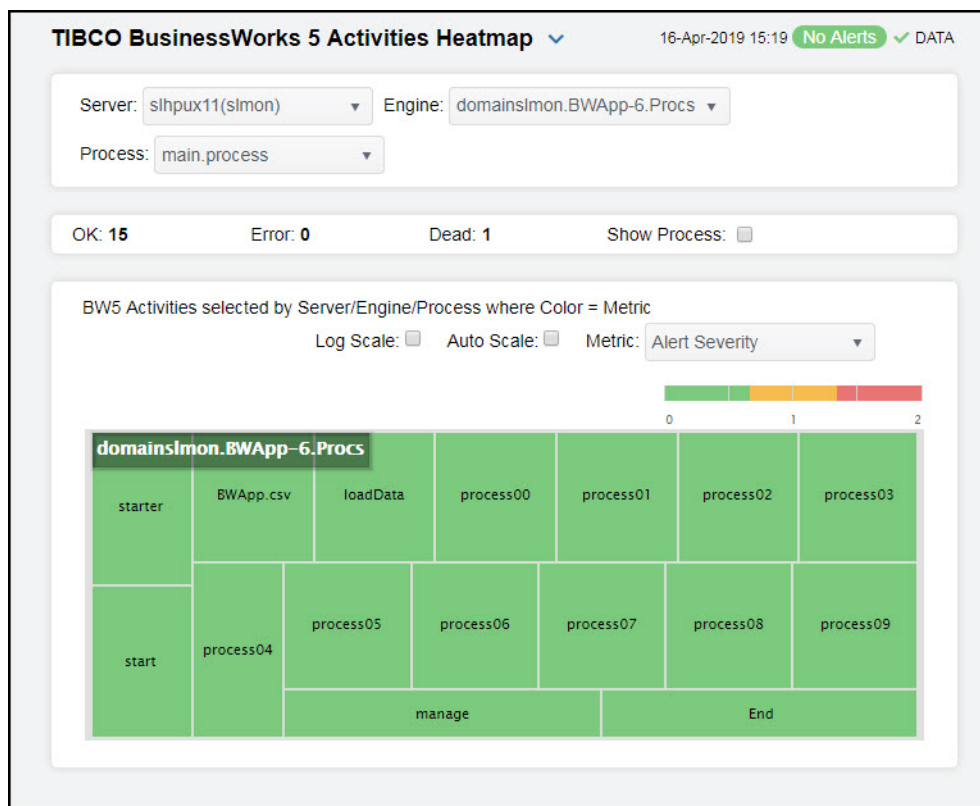
- Activity** Name of activity.
- Alert Level** The most critical alert state for alerts in the row:
 - Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
 - Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
 - Green indicates that no metrics have exceeded their alert thresholds.

Alert Count	The total number of alerts for the activity.
Time Since Last Update	Time since the last update.
Last Return Code	Status code (OK, DEAD, or ERROR) returned by most recent execution of this activity.
Exec Time ms	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.
Current Exec Time ms	Execution time (in milliseconds) accumulated during this update cycle.
Exec Time/s	Execution time accumulated per second.
Min Exec Time ms	Time (in milliseconds) of the activity that has the shortest execution time.
Max Exec Time ms	Time (in milliseconds) of the activity that has the longest execution time.
Elapsed Time ms	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.
Current Elapsed Time ms	Change in ElapsedTime this update.
Elapsed Time/s	Change in ElapsedTime per second.
Min Elapsed Time ms	Elapsed clock time (in milliseconds) of the activity that has the shortest execution time.
Max Elapsed Time ms	Elapsed clock time (in milliseconds) of the activity that has the longest execution time.
Executions	Number of times the activity has been executed.
Executions/s	Change in ExecutionCount per second.
Current Executions	Change in ExecutionCount this update.
Errors	Total number of executions of the activity that have returned an error.
Errors/s	Change in ErrorCount per second.
Current Errors	Change in ErrorCount this update.
Domain	Name of TIBCO Domain.
Activity Class	Name of the class that implements the activity.
Called Process Defs	A comma-separated list of definitions called by this activity.
Tracing	<ul style="list-style-type: none"> • true Tracing is enabled for this activity. • false Tracing is disabled for this activity.
MicroAgentInstance	Unique ID of the microagent reporting the metrics.
Executions Since Reset	Number of times the activity has been executed since the last reset of the statistics.
Source	Name of RTView Data Server sending this data (or localhost).
Time Stamp	Time of this update.

TIBCO BusinessWorks 5 Activities Heatmap

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"](#) 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.




Filter By:

The display might include these filtering options:


- Server:** Select the server for which you want to view data in the display.
- Engine:** Select the engine for which you want to view data in the display.
- 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 .
Show Process	Select this check box to display the names of the processes above their respective rectangles in the heatmap.
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.
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 Scale is not selected.
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 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 " display for a detailed view of metrics for that particular activity.
Alert Severity	<p>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:</p> <ul style="list-style-type: none">  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.
Executions	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 0 to the number of executions in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
Errors	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 number of errors in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
Exec Time/sec	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 0 to the rate of executions in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
Errors/sec	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 rate of errors in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

**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 most recent execution time in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

**Max Exec
Time**

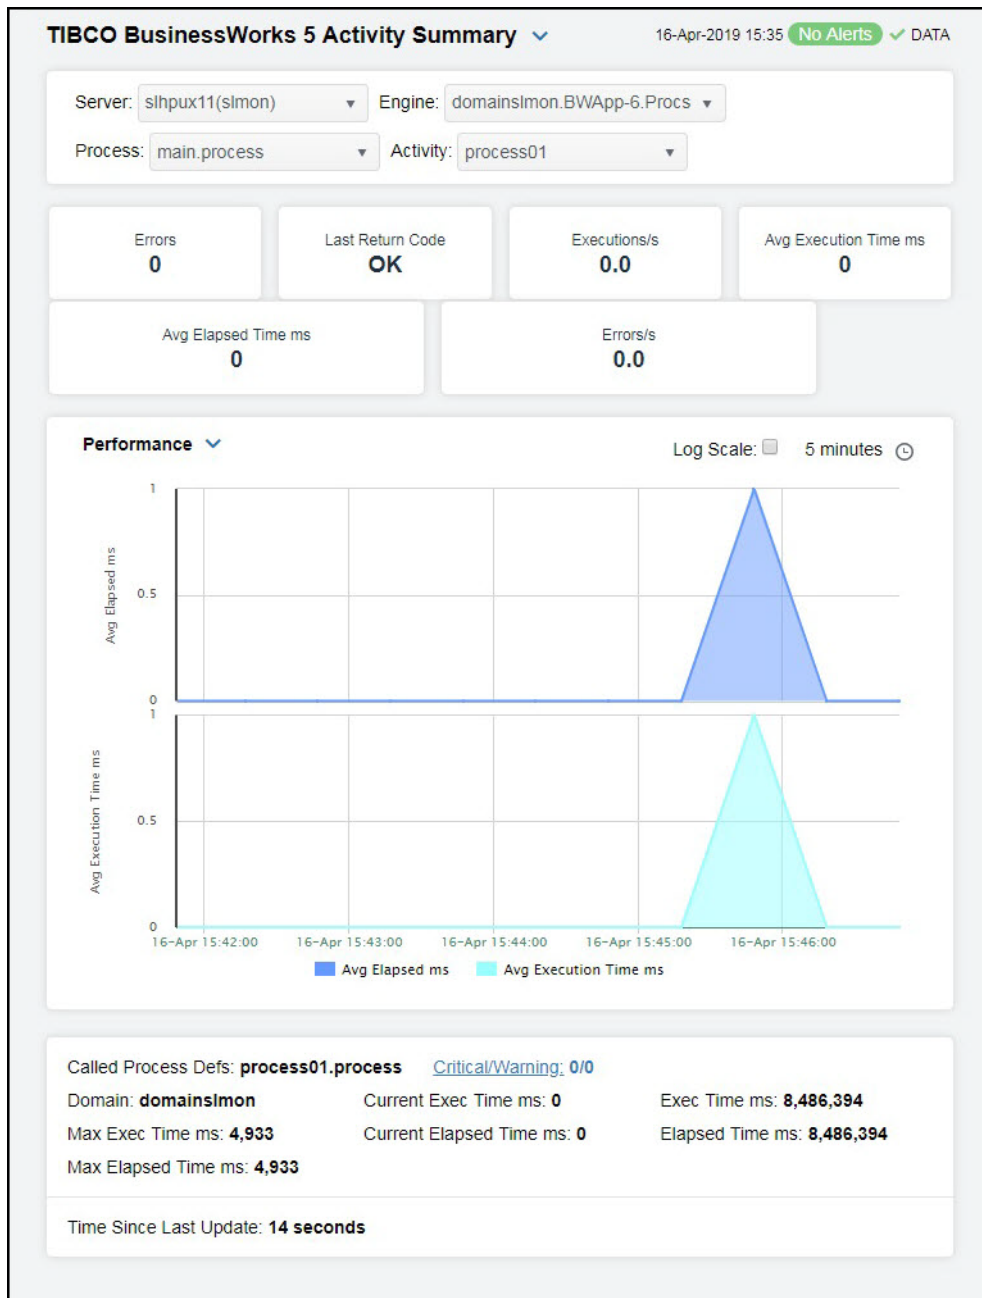
The maximum execution time for executed processes in the heatmap rectangle. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum execution time in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

TIBCO BusinessWorks 5 Activity Summary

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](#)" 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.



Filter By:

The display might include these filtering options:

- 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.
- 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.

Fields and Data

Errors	Total number of errors generated by the activity.
Last Return Code	The return code of the last execution of this activity.
Executions /s	The rate of executions, per second, for this activity.
Avg Execution Time ms	The average execution time, in milliseconds, for this activity.
Avg Elapsed Time ms	The average elapsed time, in milliseconds, for this activity.
Errors/s	The rate of errors, per second, for this activity.

Trend Graphs**Performance**

Avg Elapsed Time ms -- Traces the average elapsed time, in milliseconds, for this activity.

Avg Execution Time ms -- Traces the average execution time, in milliseconds, for this activity.

Success Rate and Avg Failures

Executions/s -- Traces the rate of executions, per second, for this activity.


Avg Errors -- Traces the average number of errors generated by the activity.

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

Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

Called Process Defs	The number of call process definitions for the activity.
Current Exec Time ms	The time spent, in milliseconds, in execution since the last data update.
Current Elapsed Time ms	The elapsed time, in milliseconds, accumulated since the last data update.
Critical/Warning	The number of critical and warning alerts.
Exec Time ms	The shortest amount of time needed to execute a process.
Elapsed Time ms	The shortest amount of elapsed time.
Domain	The name of the domain.
Max Exec Time ms	The longest amount of time needed to execute a process.
Max Elapsed Time ms	The longest amount of elapsed time.
Time Since Last Update	The amount of time since the last update.

TIBCO Enterprise Message Service

The HTML version features an overview display, "[TIBCO EMS Overview](#)" (pictured below), and the following Views which can be found under **Components tab > Middleware > TIBCO EMS**.

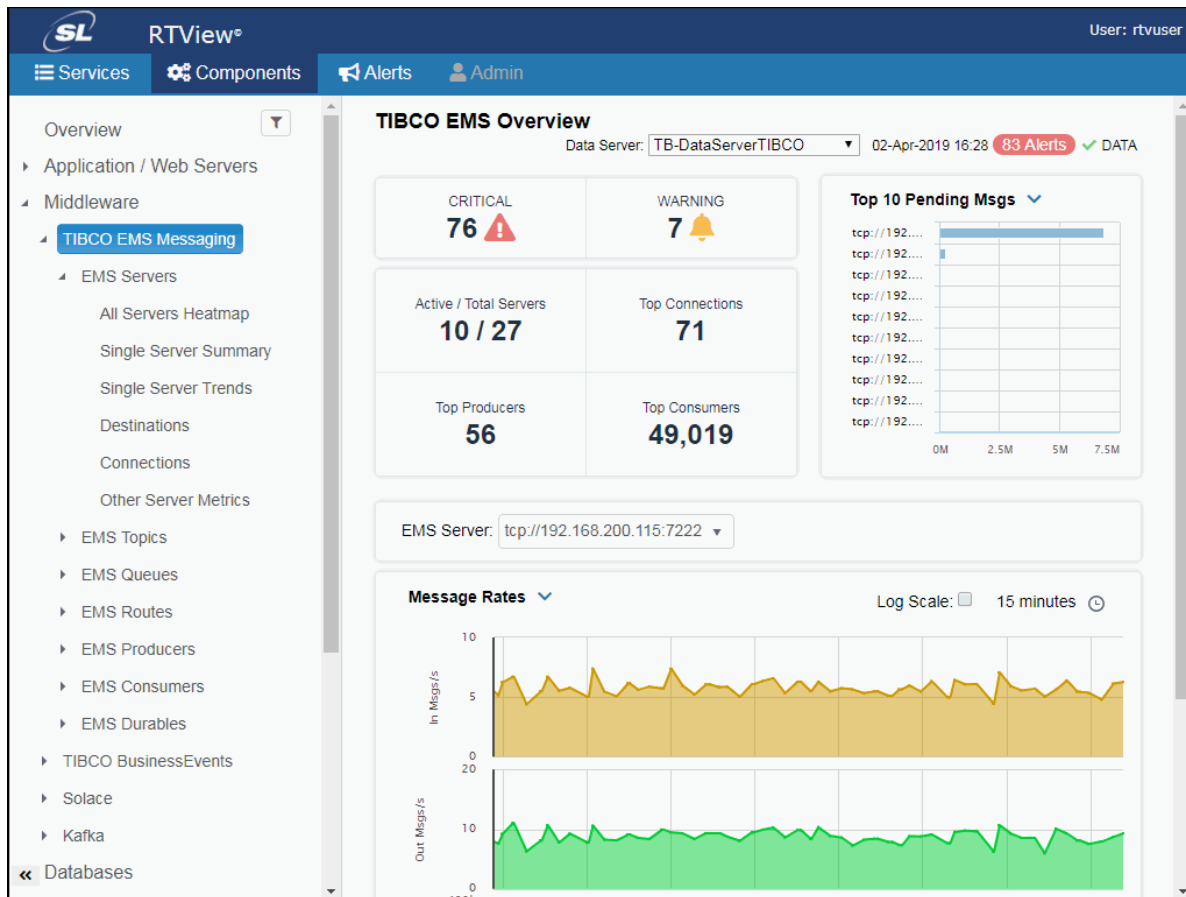
TIBCO EMS Overview

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.



This section includes descriptions of the EMS Monitor Views and their associated displays.

- “EMS Servers”
- “EMS Topics”:
- “EMS Queues”
- “EMS Routes”
- “EMS Producers”
- “EMS Consumers”
- “EMS Durables”

EMS Servers

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”](#) 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”](#), which shows server and alert status for all EMS servers.
- **Single Server Summary**: Opens the [“TIBCO EMS Server Summary”](#) 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”](#) 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”](#) display, which shows destination data for a selected server.
- **Connections**: Opens the [“TIBCO EMS Server Connections”](#) display, which shows connection information for a selected server.
- **Other Server Metrics**: Opens the [“TIBCO EMS Bridges, Users, Ports”](#) display, which shows bridges data, along with associated users and ports, for a selected server.

TIBCO EMS Servers Table

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”](#) 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	🔄
tcp://192.168.200.138:7222	Unknown (tcp://192.168.200.138:7222)	192.168.200.138	🔄
tcp://192.168.200.34:7222	EMS-SERVER-TB34	192.168.200.34	
tcp://SLHOST10:7010	Unknown (tcp://SLHOST10:7010)	SLHOST10	🔄
tcp://SLHOST10:7011	Unknown (tcp://SLHOST10:7011)	SLHOST10	🔄
tcp://SLHOST10:7021	Unknown (tcp://SLHOST10:7021)	SLHOST10	🔄
tcp://SLHOST10:7020	Unknown (tcp://SLHOST10:7020)	SLHOST10	🔄
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	

Fields and Data

This display includes:

- Servers** 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.
- Max In Msgs/s** The highest rate of inbound messages, per second, from all producers and consumers on all EMS servers.
- Max Out Msgs/s** The highest rate of outbound messages, per second, from all producers and consumers on all EMS servers.
- Total Pending Msgs** The total number of inbound and outbound messages waiting to be processed on all EMS servers.
- Active Only** Select this check box to display only the active servers in the table below.
- Table** This table shows information for all EMS servers. Click on a table row to drill-down to the ["TIBCO EMS Server Summary"](#) display and view metrics for that particular server.
- URL** Select to include servers that are currently in Standby mode. **Standby Servers** are represented in blue.
- Server Name** The name of the server.
- Host** The name or IP address for the host server.

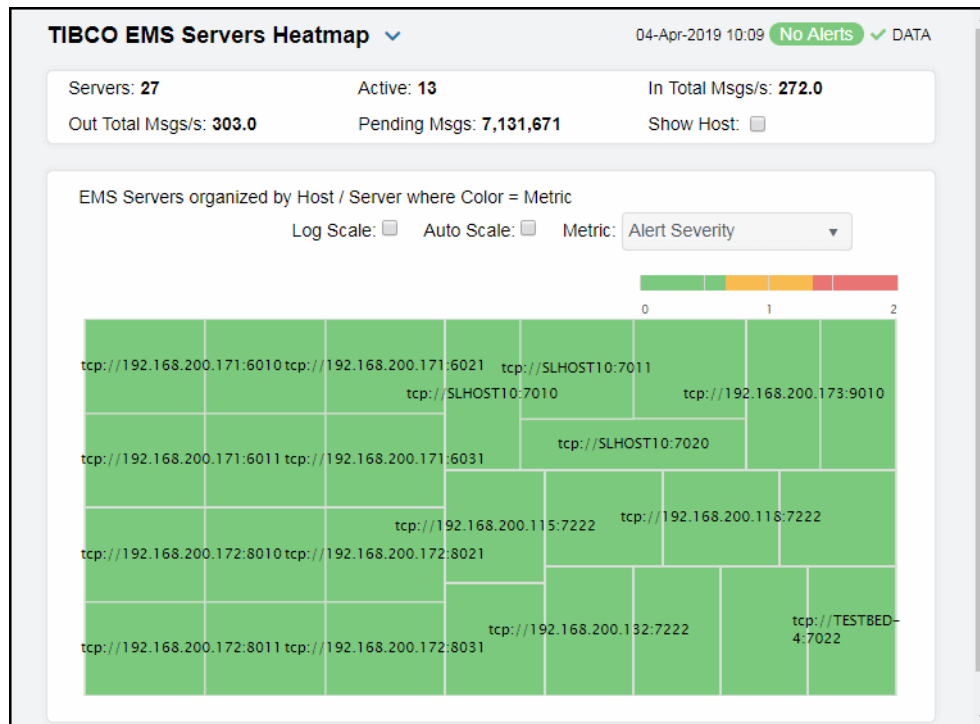
Alert Level	<p>The current alert level.</p> <ul style="list-style-type: none">  -- One or more alerts have exceeded their specified ALARMLEVEL threshold, have an Alert Severity value of 2, and are shown in red.  -- One or more alerts have exceeded their specified WARNINGLEVEL threshold, have an Alert 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 number of current alerts.
State	<p>The server status:</p> <p>Active -- The server is currently processing requests.</p> <p>Inactive -- The server is not currently processing requests. Inactive Servers are represented in dark red.</p> <p>Standby -- The server is functioning as a backup for a primary server. Standby Servers are represented in blue.</p>
Pending Msgs	The number of currently pending messages on the server.
Connections	The number of clients currently connected to the server.
Disk Reads/s	The speed at which the server reads disk data.
Disk Writes/s	The speed at which the server writes data to disk.
Durables	The number of durables on the server.
Msgs In	The number of inbound messages received by the server since the server was started.
In Msgs/s	The rate of inbound messages in number of messages per second.
Max Msg Memory Bytes	The maximum amount of memory, in bytes, allocated for use by messages on the server.
Msg Memory Bytes	The amount of memory, in bytes, currently used by messages on the server.
Msg Memory %	The amount of memory, in percent, used by messages on the server.
Msg Mem Pooled	The currently allocated pool size, in bytes, for messages.
Out Msgs	The number of outbound messages sent by the server since the server was started.
Out Msgs/s	The rate of outbound messages in number of messages per second.
Pending Msg Size	The amount of space, in bytes, pending messages use on the server.
Version	The TIBCO EMS software version currently running.
FT URL	The IP address and port number for the source (application, server, and so forth) associated with the alert.
Async DB Size Bytes	The amount of database space, in bytes, occupied by asynchronous data on the server.

Backup Name	The name of the backup server assigned as the backup to this server.
PID	The process ID of the EMS server.
Queues	The number of message queues.
Start Time	The date and time that the server was started.
Sync DB Size Bytes	The amount of database space, in bytes, occupied by synchronous data on the server.
Topics	The number of currently active topics on the server.
Uptime	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.

TIBCO EMS Servers Heatmap

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 (%)**.

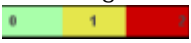



The heatmap is organized by host, each rectangle representing a server. The rectangle color indicates the most critical alert state. Click on a node to drill-down to the "[TIBCO EMS Server Summary](#)" display and view metrics for a particular server. Toggle between the commonly accessed **Table** and **Heatmap** displays by clicking the drop down list on the display title. Mouse-over rectangles to view more details about host performance and status.






Fields and Data

This display includes:

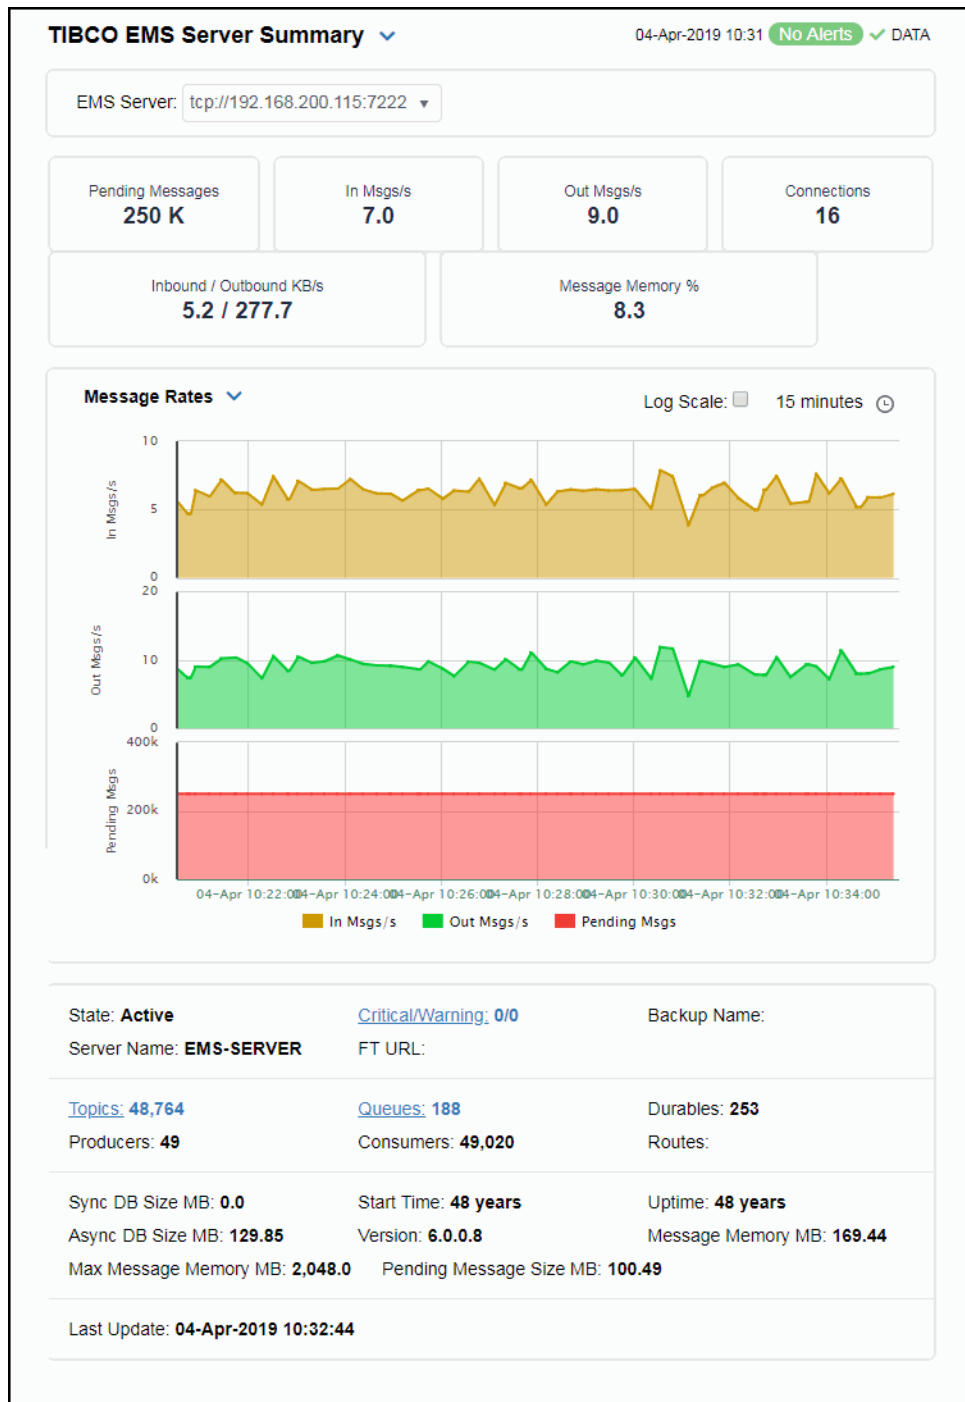
Servers	The total number of active, inactive, and standby EMS servers.
Active	The total number of currently active EMS servers.
In Total Msgs/s	The total number of inbound messages, per second, from all producers and consumers on all EMS servers.
Out Total Msgs/s	The total number of outbound messages, per second, from all producers and consumers on all EMS servers.
Pending Msgs	The total number of pending messages waiting to be processed on all EMS servers. Click to open the "TIBCO EMS Servers Table" display.
Show Host	Select this check box to display the name of the host for 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.

Auto Scale	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 "TIBCO EMS Server Summary" display for a detailed view of metrics for that particular server.
Alert Severity	<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 2 is the greatest Alert Severity.</p> <p>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.</p> <p>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.</p> <p>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.</p>
Alert Count	<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 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>
Connections	<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 0 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 Auto option does not impact this metric.</p>
Pending Msgs	<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 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).</p> <p>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.</p>

In Msgs/s	<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 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).</p> <p>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.</p>
Out Msgs/s	<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 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).</p> <p>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.</p>
Msg Mem %	<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 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).</p> <p>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.</p>

TIBCO EMS Server Summary

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](#)" 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](#)" display. Clicking the **Queues** link at the bottom of the display opens the TIBCO EMS Queues Table display.



Fields and Data

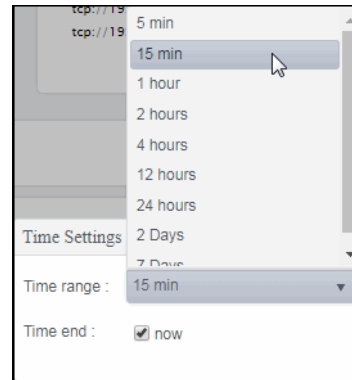
This display includes:

- EMS Server** Select the EMS Server for which you want to view data. The selection made here populates this display.

Pending Messages	The total number of pending messages on the server.
In Msgs/s	The rate of inbound messages on the server.
Out Msgs/s	The rate of outbound messages on the server.
Connections	The number of connections on the server.
Inbound/Outbound KB/s	The rate of inbound and outbound kilobytes on the server.
Message Memory %	The percentage of message memory utilization on the server.
Trend Graphs	<p>Message Rates Trend Graph</p> <p>In Msgs / sec -- Traces the number of inbound messages, per second, from all producers and consumers.</p> <p>Out Msgs / sec -- Traces the number of outbound messages, per second, from all producers and consumers.</p> <p>Pending Msgs -- Traces the total number of inbound and outbound messages currently waiting to be processed.</p> <p>Message Flows Trend Graph</p> <p>In Msgs -- Traces the total number of inbound messages from all producers and consumers.</p> <p>Out Msgs -- Traces the number of outbound messages from all producers and consumers.</p> <p>Pending Msgs -- Traces the total number of inbound and outbound messages currently waiting to be processed.</p> <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.</p>

Time Settings

Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

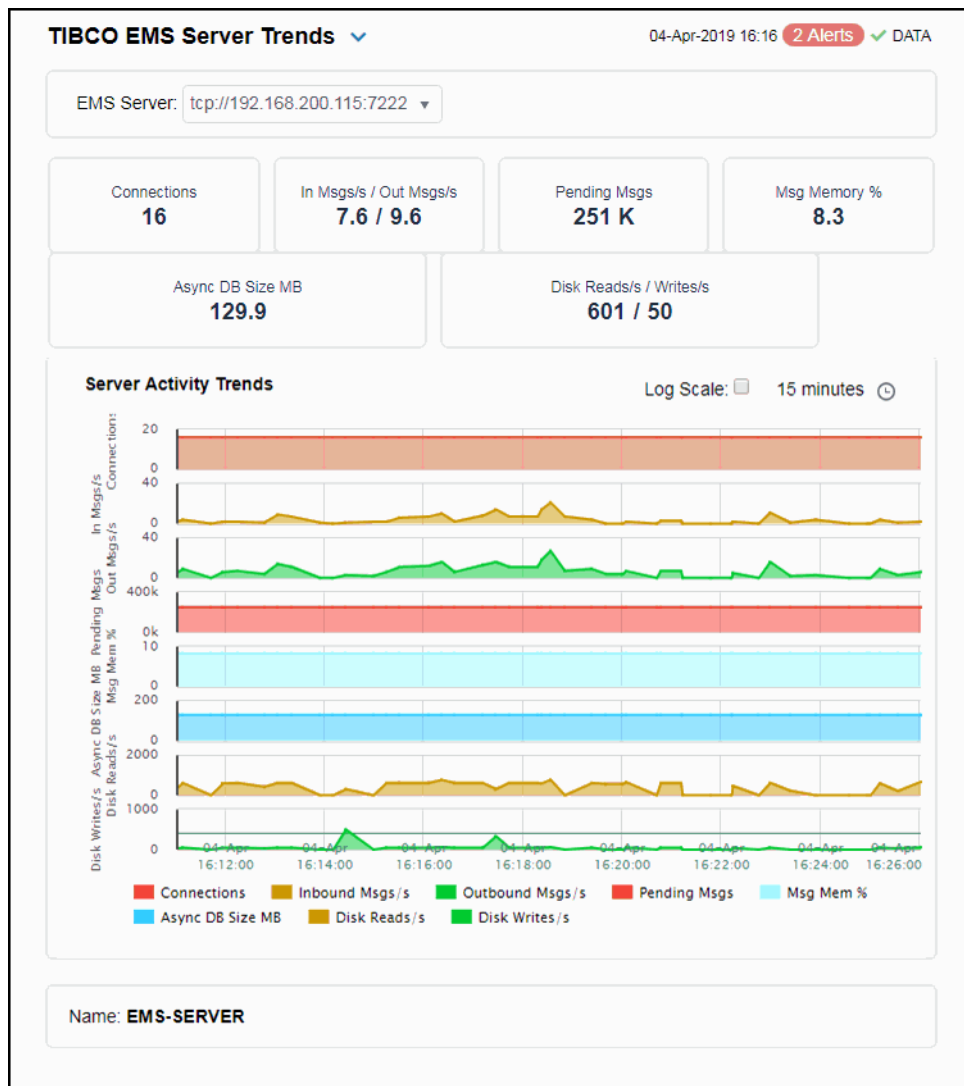
Server Information

State	The server status: Active -- The server is currently processing requests. Inactive -- The server is not currently processing requests. Standby -- The server is functioning as a backup for a primary server.
Critical/Warning	Lists the number of critical and warning level alerts on the server.
Backup	The name of the backup server for the server.
Server Name	The name of the selected server.
FT URL	The IP address and port number, or the hostname and port number, of the fault tolerant standby server assigned to this server.
Topics	The number of topics currently active on the server. Click to open the "TIBCO EMS Topics Table" display for details.
Queues	The number of queues currently active on the server. Click to open the "TIBCO EMS Queues Table" display for details.
Durables	The number of durables currently active on the server.
Producers	The number of producers currently active on the server.

Consumers	The number of consumers currently connected to the server.
Routes	The number of routes defined on the server.
Sync DB Size MB	The amount of database space, in megabytes, used by synchronous message persistence data on the server.
Start Time	The date and time that the server was started.
Uptime	The amount of time since the server was started. Format: dd HH:MM:SS <days> <hours>:<minutes>:<seconds> For example: 10d 08:41:38
Async DB Size MB	The amount of database space, in megabytes, used by asynchronous message persistence data on the server.
Version	The TIBCO EMS software version currently running.
Message Memory MB	The amount of memory, in megabytes, used by message persistence on the server.
Max Message Memory MB	The maximum amount of memory, in megabytes, used by message persistence on the server.
Pending Message Size MB	The total size of inbound and outbound messages, in megabytes, currently waiting to be processed.
Last Update	The time that a data update was last made.

TIBCO EMS Server Trends

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"](#) 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.



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.
- Connections** The total number of client connections.

In Msgs/s / Out Msgs/s	The number of inbound messages, per second, from all producers and consumers and the number of outbound messages, per second, from all producers and consumers.
Pending Msgs	The total number of messages currently waiting to be processed.
Msg Memory %	The amount of memory, in percent, used by messages.
Async DB Size MB	The amount of database space, in megabytes, used by asynchronous data on the server.
Disk Reads/s / Writes/s	The amount of disk data, in kilobytes, read by the server since the server was started and the amount of data, in kilobytes, written to disk by the server since the server was started.
Server Activity Trends	Shows metrics for the selected server. <ul style="list-style-type: none"> 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.

Time Settings

Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

- Server Name** The name of the EMS Server selected from the EMS Server drop-down menu.
- Critical/Warning** The total number of critical and warning alerts for the server.
- Last Update** The date and time of the last data update.

TIBCO EMS Server Destinations

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**

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.
- Pending Messages** The total number of messages currently waiting to be processed.
- In Msgs/s** The number of inbound messages, per second, from all producers and consumers.

Out Msgs/s	The number of outbound messages, per second, from all producers and consumers.																																						
Connections	The total number of client connections.																																						
In/Out KB/s	The rate of incoming kilobytes (per second) and the rate of outgoing kilobytes (per second).																																						
Message Memory %	The amount of memory, in percent, used by messages.																																						
Queues Table	<table> <tr> <td>Queue Name</td> <td>The name of the queue.</td> </tr> <tr> <td>In Msgs/s</td> <td>The number of inbound messages for the queue, per second. This metric comes directly from the tibjms.admin.DestinationInfo class from TIBCO.</td> </tr> <tr> <td>In Total Msgs</td> <td>The total number of inbound messages for the queue.</td> </tr> <tr> <td>Out Msgs/s</td> <td>The number of outbound messages for the queue, per second. This metric comes directly from the tibjms.admin.DestinationInfo class from TIBCO.</td> </tr> <tr> <td>Out Total Msgs</td> <td>The total number of outbound messages for the queue.</td> </tr> <tr> <td>Pending Msgs</td> <td>The number of currently pending messages for the queue.</td> </tr> <tr> <td>Pending Msg Size</td> <td>The amount of space, in bytes, used by pending messages for the queue.</td> </tr> <tr> <td>Consumers</td> <td>The number of active and inactive consumers.</td> </tr> <tr> <td>Fail Safe</td> <td>When checked, the message is marked as failsafe delivery.</td> </tr> <tr> <td>Flow Control Max Bytes</td> <td>The maximum number of bytes allocated for use by flow control.</td> </tr> <tr> <td>Global</td> <td>When checked, the message is global and is routed to other servers.</td> </tr> <tr> <td>In KB/s</td> <td>The amount of inbound messages for the queue, in kilobytes per second.</td> </tr> <tr> <td>In KB</td> <td>The total amount of inbound messages for the queue, in kilobytes.</td> </tr> <tr> <td>Max Bytes</td> <td>The maximum amount of bytes allocated for use by the queue.</td> </tr> <tr> <td>Max Msgs</td> <td>The maximum number of messages allocated for use by the queue.</td> </tr> <tr> <td>Overflow Policy</td> <td>Indicates whether an overflow policy is set for the queue: 0 = No policy is set. 1 = A policy is set.</td> </tr> <tr> <td>Secure</td> <td>When checked, the queue is designated as secure and enforces permission policies.</td> </tr> <tr> <td>Static</td> <td>When checked, the queue has a static destination.</td> </tr> <tr> <td>Description</td> <td>Descriptive text to help the administrator identify this resource.</td> </tr> </table>	Queue Name	The name of the queue.	In Msgs/s	The number of inbound messages for the queue, per second. This metric comes directly from the tibjms.admin.DestinationInfo class from TIBCO.	In Total Msgs	The total number of inbound messages for the queue.	Out Msgs/s	The number of outbound messages for the queue, per second. This metric comes directly from the tibjms.admin.DestinationInfo class from TIBCO.	Out Total Msgs	The total number of outbound messages for the queue.	Pending Msgs	The number of currently pending messages for the queue.	Pending Msg Size	The amount of space, in bytes, used by pending messages for the queue.	Consumers	The number of active and inactive consumers.	Fail Safe	When checked, the message is marked as failsafe delivery.	Flow Control Max Bytes	The maximum number of bytes allocated for use by flow control.	Global	When checked, the message is global and is routed to other servers.	In KB/s	The amount of inbound messages for the queue, in kilobytes per second.	In KB	The total amount of inbound messages for the queue, in kilobytes.	Max Bytes	The maximum amount of bytes allocated for use by the queue.	Max Msgs	The maximum number of messages allocated for use by the queue.	Overflow Policy	Indicates whether an overflow policy is set for the queue: 0 = No policy is set. 1 = A policy is set.	Secure	When checked, the queue is designated as secure and enforces permission policies.	Static	When checked, the queue has a static destination.	Description	Descriptive text to help the administrator identify this resource.
Queue Name	The name of the queue.																																						
In Msgs/s	The number of inbound messages for the queue, per second. This metric comes directly from the tibjms.admin.DestinationInfo class from TIBCO.																																						
In Total Msgs	The total number of inbound messages for the queue.																																						
Out Msgs/s	The number of outbound messages for the queue, per second. This metric comes directly from the tibjms.admin.DestinationInfo class from TIBCO.																																						
Out Total Msgs	The total number of outbound messages for the queue.																																						
Pending Msgs	The number of currently pending messages for the queue.																																						
Pending Msg Size	The amount of space, in bytes, used by pending messages for the queue.																																						
Consumers	The number of active and inactive consumers.																																						
Fail Safe	When checked, the message is marked as failsafe delivery.																																						
Flow Control Max Bytes	The maximum number of bytes allocated for use by flow control.																																						
Global	When checked, the message is global and is routed to other servers.																																						
In KB/s	The amount of inbound messages for the queue, in kilobytes per second.																																						
In KB	The total amount of inbound messages for the queue, in kilobytes.																																						
Max Bytes	The maximum amount of bytes allocated for use by the queue.																																						
Max Msgs	The maximum number of messages allocated for use by the queue.																																						
Overflow Policy	Indicates whether an overflow policy is set for the queue: 0 = No policy is set. 1 = A policy is set.																																						
Secure	When checked, the queue is designated as secure and enforces permission policies.																																						
Static	When checked, the queue has a static destination.																																						
Description	Descriptive text to help the administrator identify this resource.																																						

	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.
	Out KB/s	The amount of outbound messages for the queue, in kilobytes per second.
	Out KB	The total amount of outbound messages for the queue, in kilobytes.
	Exclusive	When checked, the server sends all messages on this queue to one consumer.
	Max Redelivery	The maximum number of attempts for attempting redelivery of a message.
	Filter In Patterns	The string used to filter the data in the row.
	Receivers	The number of receivers that receive queue message.
Topics Table	Topic Name	The name of the topic.
	In Msgs/s	The number of inbound messages for the topic, per second. Note: This metric comes directly from the tibjms.admin.DestinationInfo class from TIBCO.
	In Total Msgs	The total number of inbound messages for the topic.
	Out Msgs/s	The number of outbound messages for the topic, per second. Note: This metric comes directly from the tibjms.admin.DestinationInfo class from TIBCO.
	Out Total Msgs	The total number of outbound messages for the topic.
	Pending Msgs	The number of currently pending messages for the topic.
	Pending Msg Size	The amount of space, in bytes, used by pending messages for the topic.
	Active Durables	The number of currently active durables or the topic.
	Consumers	The number of consumers for the topic.
	Durables	The number of durables for the topic.
	Fail Safe	When checked, the message is marked as failsafe delivery.
	Flow Control Max Bytes	The maximum number of bytes allocated for use by flow control.
	Global	When checked, the message is global and is routed to other servers.
	In KB/s	The amount of inbound messages for the topic, in kilobytes per second.

In KB	The total amount of inbound messages for the topic, in kilobytes, since the server started.
Max Bytes	The maximum size, in bytes, that the topic can store for delivery to each durable or non-durable online subscriber on that topic.
Max Msgs	The maximum number of messages before the server indicates an error and overflow policies are activated.
Out KB/s	The amount of outbound messages for the topic, in kilobytes per second.
Out KB	The total amount of outbound messages for the topic, in bytes.
Overflow Policy	Indicates whether an overflow policy is set for the topic: 0 = No policy is set. 1 = A policy is set.
Secure	When checked, the topic is designated as secure and enforces permission policies.
Static	When checked, the topic has a static destination.
Subscribers	The number of subscribers for the topic.
Description	Descriptive text to help the administrator identify this resource.
Current in Total Msgs	Displays the change (delta) in inboundTotalMessages from the previous cache refresh to the current cache refresh.
Current in Total Bytes	Displays the change (delta) in inboundTotalBytes from the previous cache refresh to the current cache refresh.
Current Out Total Msgs	Displays the change (delta) in outboundTotalMessages from the previous cache refresh to the current cache refresh.
Current Out Total Bytes	Displays the change (delta) in outboundTotalBytes from the previous cache refresh to the current cache refresh.
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.
Server Name	The name of the EMS Server selected from the EMS Server drop-down menu.
Critical/Warning	The total number of critical and warning alerts for the server.
Last Update	The date and time of the last data update.

TIBCO EMS Server Connections

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](#)” 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

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.
Connections Table	This table describes the current connections on the selected server.
Conn ID	The unique numeric ID assigned to this connection that can be used for deletion.
Client ID	The unique string identifier assigned to the client.
connectionURL	The connection URL.
User Name	The user name.

Host	The name of the host to which the server is connected.
Type	The type of connection: Queue, Topic or System.
Consumers	The total number of consumers currently connected.
Producers	The total number of producers currently connected.
Sessions	The total number of sessions currently connected.
Start Time	The date and time the server was started
Up Time	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.
State	The server status: Active -- The server is currently processing requests. Inactive -- The server is not currently processing requests. Standby -- The server is functioning as a backup for a primary server.
Critical/Warning	The total number of critical and warning alerts for the server.
Backup Name	The name of the backup server for this server.
Server Name	The name of the EMS Server selected from the EMS Server drop-down menu.
FT URL	The IP address and port number, or the hostname and port number, of the fault tolerant standby server assigned to this server.
Topics	The number of topics currently active on the server. Click to open the "TIBCO EMS Topics Table" display for details.
Queues	The number of queues currently active on the server. Click to open the "TIBCO EMS Queues Table" display for details.
Durables	The number of durables currently active on the server.
Producers	The number of producers currently active on the server.
Consumers	The number of consumers currently active on the server.
Routes	The number of routes defined on the server.
Last Update	The date and time of the last data update.

TIBCO EMS Bridges, Users, Ports

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: tcp://192.168.200.115:7222

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**

Fields and Data

This display includes:

- EMS Server** The EMS Server selected from this drop-down menu populates all associated Bridges, Users, and Ports data in this display.
- Bridges** This table describes the bridges for the selected server.

source	The topic or queue which is the source of the bridge.
target	The topic or queue which is the target of the bridge.
Selector	The message selector string or blank if none has been set.
Source Type	The type of the source.
Target Type	The type of the target.
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.
Users	This table describes the users on the selected server.
Description	Textual description of the user.
External	When checked, the user is defined in an external system.
Name	The name of the connected user.
Listen Ports	This table describes the connections the selected server is to listen for.
port	The IP address and port number on which the server is to listen for connections.
URL	The URL on which the server is to listen for connections.
Name	The name of the EMS Server selected from the EMS Server drop-down menu.
Connections	The number of connections on the server.

EMS Topics

These displays present several views of performance metrics for topics. Clicking **EMS Topics** from the left/navigation menu opens the [“TIBCO EMS Topics Table”](#) 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”](#) 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”](#) display. The options available under **EMS Topics** are:

- **All Topics Summary:** Clicking **All Topics Summary** opens the [“TIBCO EMS Topics for Server Summary”](#) 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”](#), 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”](#) 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”](#), 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

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"** 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 1 of 8 1 - 40 of 306 items

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 by Topic Name 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.




Shown The total number of currently active topics on the selected server, which is filtered by the **Data Collection > Metric Filters > Topics** field in the RTView Configuration Application. The default value for the **Topics** property is:

```
^(?!^\\$sys\\.|^\\$TMP\\.|^AMX_MGMT\\.|^EMSGMS\\.|^AMX_SV\\.|^_HAWK\\.|^_LOCAL\\.|^_HAWK\\.|^TMP\\.EMS)
```

You can modify the filter value by editing the **Topics** property, which will override the default value. See [“Configuring Data Collection”](#) for more information.

Total The total number of topics on the selected server.

Table This table describes all topics on the selected server. Click a row to view metrics for a single topic in the [“TIBCO EMS Topic Summary”](#) display.

Topic Name	The name of the topic.
URL	The IP address and port number for the server.
Alert Level	The current alert level.  -- 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 number of current alerts.
In Msgs/s	The number of inbound messages for the topic, per second.
In Total Msgs	The total number of inbound messages for the topic.
Out Msgs/s	The number of outbound messages for the topic, per second.
Out Total Messages	The total number of outbound messages for the topic.
Pending Msgs	The number of currently pending messages for the topic.
Pending Msg Size	The amount of space, in bytes, used by pending messages for the topic.
Active Durables	The number of currently active durables or the topic.
Consumers	The number of consumers for the topic.
Durables	The number of durables for the topic.
Fail Safe	When checked, the message is marked as failsafe delivery.
Flow Control Max Bytes	The maximum number of bytes allocated for use by flow control.
Global	When checked, the message is global and is routed to other servers.
In KB/s	The amount of inbound messages for the topic, in kilobytes per second.

In MB	The total amount of inbound messages for the topic, in megabytes, since the server started.
Max Bytes	The maximum size, in bytes, that the topic can store for delivery to each durable or non-durable online subscriber on that topic.
Max Msgs	The maximum number of messages before the server indicates an error and overflow policies are activated.
Out KB/s	The amount of outbound messages for the topic, in kilobytes per second.
Out MB	The total amount of outbound messages for the topic, in megabytes.
Overflow Policy	Indicates whether an overflow policy is set for the topic: 0 = No policy is set. 1 = A policy is set.
Secure	When checked, the topic is designated as secure and enforces permission policies.
Static	When checked, the topic has a static destination.
Subscribers	The number of subscribers for the topic.
Description	Descriptive text to help the administrator identify this resource.
Current In Total Messages	Displays the change (delta) in <code>inboundTotalMessages</code> from the previous cache refresh to the current cache refresh.
Current In Total Bytes	Displays the change (delta) in <code>inboundTotalBytes</code> from the previous cache refresh to the current cache refresh.
Current Out Total Msgs	Displays the change (delta) in <code>outboundTotalMessages</code> from the previous cache refresh to the current cache refresh.
Current Out Total Bytes	Displays the change (delta) in <code>outboundTotalBytes</code> from the previous cache refresh to the current cache refresh.
In Msgs/s (TIBCO)	The number of inbound TIBCO messages for the topic, per second. This metric comes directly from the <code>tibjms.admin.DestinationInfo</code> class from TIBCO.
Out Msgs/s (TIBCO)	The number of outbound TIBCO messages for the topic, per second. This metric comes directly from the <code>tibjms.admin.DestinationInfo</code> class from TIBCO.
In KB/s (TIBCO)	The amount of inbound TIBCO messages for the topic, in kilobytes per second.
Out KB/s (TIBCO)	The amount of outbound TIBCO messages for the topic, in kilobytes per second.

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.

TIBCO EMS Topics for Server Summary

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"** 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.




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.
Consumers	The number of consumers currently active on the server.
Subscribers	The number of subscribers for the topic.
In Msgs/s	The number of inbound messages for the topic, per second.
Out Msgs/s	The number of outbound messages for the topic, per second.
Pending Messages	The number of currently pending messages for the topic.
Pending Msg Size MB	The amount of space, in bytes, used by pending messages for the topic.
Trend Graphs	The following trend graphs are available.
Message Rates	<p>In Msgs/s -- The number of inbound messages for all topics on the server, per second.</p> <p>Out Msgs/s -- The number of outbound messages for all topics on the server, per second.</p> <p>Pending Msgs -- The total number of messages for all topics on the server currently waiting to be processed.</p>
KB Rates	<p>In KB/s -- The size of inbound messages, in kilobytes per second, for all topics on the server.</p> <p>Out KB/s -- The size of outbound messages, in kilobytes per second, for all topics on the server.</p> <p>Pending Msg Size KB -- The total size of messages, in bytes, for all topics on the server currently waiting to be processed.</p>
Pending Msgs	<p>Pending Messages -- The total number of messages for all topics on the server currently waiting to be processed.</p> <p>Pending Msg Size KB-- The total size of messages, in bytes, for all topics on the server currently waiting to be processed.</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.

Time Settings

Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.




To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

Name	The name of the server selected in the EMS Server drop down list.
Durables	The number of durables for the topic.
In KB/s	The size of inbound messages, in kilobytes per second, for all topics on the server.
Out KB/s	The size of outbound messages, in kilobytes per second, for all topics on the server.

TIBCO EMS Topics Heatmap

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"](#) display and view metrics for that particular Topic.









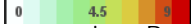

Fields and Data

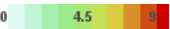

This display includes:

- EMS Server** The EMS Server selected from this drop-down menu populates all associated Topic data in this display.

Filter by Topic Name	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.
Shown	The total number of currently active topics on the selected server, which is filtered by the Data Collection > Metric Filters > Topics field in the RTView Configuration Application. The default value for the Topics property is: <code>^(?!^\\\$sys\\. ^\\\$TMP\\. ^AMX_MGMT\\. ^EMSGMS\\. ^AMX_SV\\. ^_HAWK\\. ^_LOCAL\\. _HAWK\\. ^TMP\\.EMS)</code> You can modify the filter value by editing the Topics property, which will override the default value. See "Configuring Data Collection" for more information.
Total	The total number of topics on the selected server.
Show URL	Select this check box to display the server URL 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.
Auto Scale	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 "TIBCO EMS Topic Summary" display for a detailed view of metrics for that particular topic.

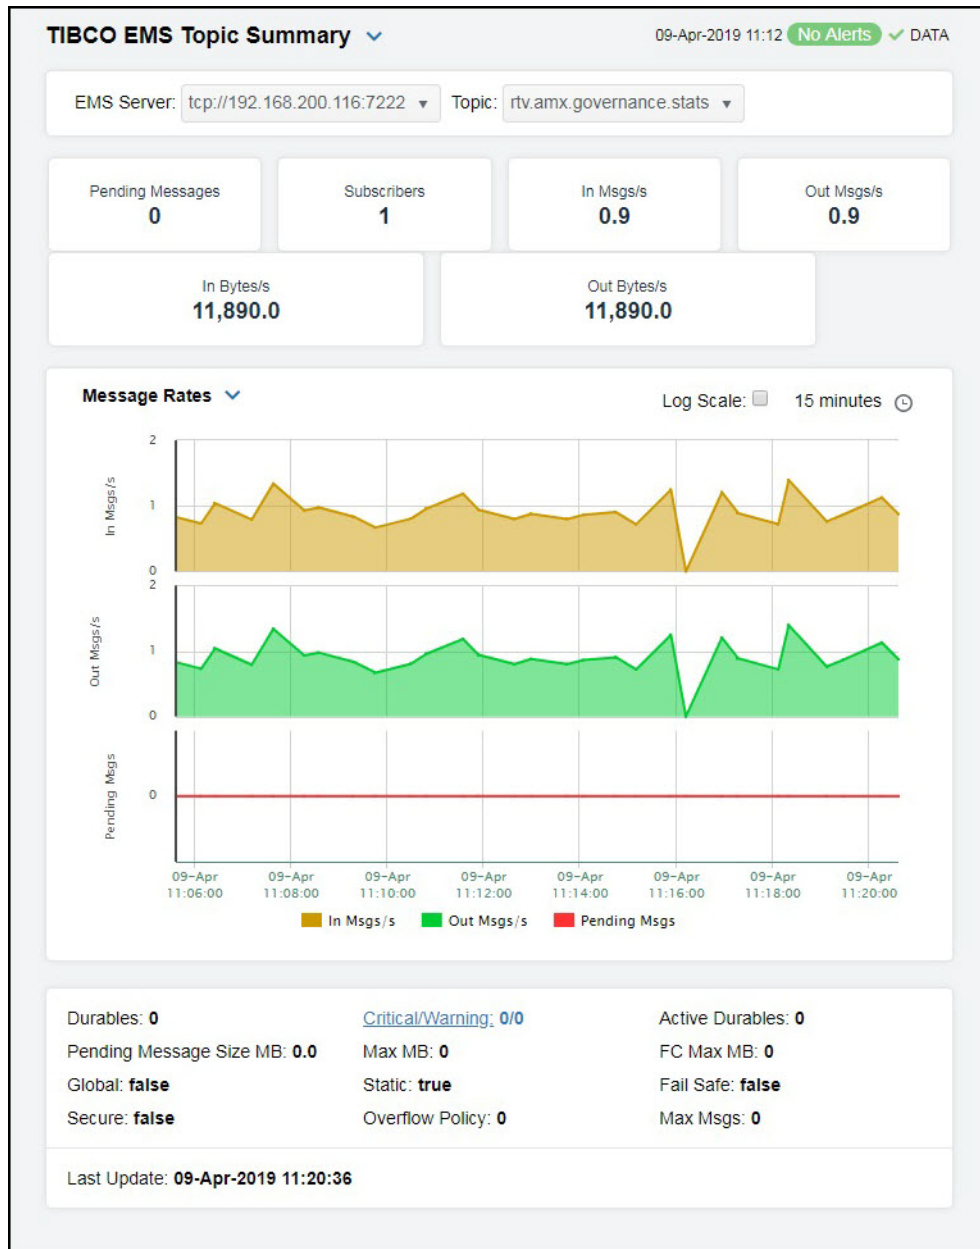
Alert Severity	<p>The maximum alert level in the item (index) associated with the rectangle. Values range from 0 to 2, as indicated in the color gradient bar , where 2 is the greatest Alert Severity.</p> <p>2 -- Metrics that have exceeded their specified ALARMLEVEL threshold and have an Alert Severity value of 2 are shown in red. For a given rectangle, this indicates that one or more metrics have exceeded their alarm threshold.</p> <p>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.</p> <p>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.</p>
Alert Count	<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 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>
Consumers	<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 0 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 Auto Scale option does not impact this metric.</p>
Durables	<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 0 to the maximum count of durables in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>
Subscribers	<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 0 to the maximum count of subscribers in the heatmap. The middle value in the gradient bar indicates the middle value of the range.</p>

- 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 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.
- 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 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.
- 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 Scale** 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 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

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"](#) 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.



Filters

- EMS Server** The EMS Server selected from this drop-down menu populates the **Topic** drop-down menu with the Topics belonging to this EMS Server.
- Topic** Select the topic for which you want to view data in the display.

Fields and Data

This display includes:

- Pending Messages** The number of messages for the selected topic currently waiting to be processed.
- Subscribers** The number of subscribers for the topic.
- In Msgs/s** The number of inbound messages, per second, for the selected topic.
- Out Msgs/s** The number of outbound messages, per second, for the selected topic.
- In KB/s** The size of inbound messages, in kilobytes per second, for the selected topic.
- Out KB/s** The size of outbound messages, in kilobytes per second, for the selected topic.

Trend Graphs**Message Rates**

In Msgs / sec -- Traces the number of inbound messages, per second. This trend graph only displays when **Use Rates** is selected.

Out Msgs / sec -- Traces the number of outbound messages, per second. This trend graph only displays when **Use Rates** is selected.

Pending Msgs -- Traces the number of messages currently waiting to be processed.

Message Flows

In Msgs -- Traces the number of inbound messages.

Out Msgs -- Traces the number of outbound messages


Pending Msgs -- Traces the number of messages currently waiting to be processed.

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.

Time Settings

Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

Durables	The number of durable subscribers (active and inactive) to the topic.
Critical/Warning	The total number of critical and warning alerts for the server.
Pending Message Size MB	The size of the messages for the selected topic, in megabytes, currently waiting to be processed.
Max MB	The maximum of memory, in megabytes, allocated for use by the topic.
FC Max MB	The maximum amount of memory, in megabytes, allocated for flow control use by the topic.
Global	When true, the message is global and is routed to other servers.
Static	When true, the topic has a static destination.
Fail Safe	When true, the message is marked as failsafe delivery.
Secure	When true, 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.

Max Msgs The maximum number of messages allocated for the topic.

Last Update The date and time of the last data update.

TIBCO EMS Topic Detail by Server

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”](#) display, where you can view additional data for that particular topic on that particular server.

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	

Filters

Topic The Topic selected from this drop-down menu populates this display.

Fields and Data

This display includes:

Count The number of topics listed in the table.

Table Shows details about the selected Topic for each server that has the Topic defined. Double-click on a table row to view details in the [“TIBCO EMS Topic Summary”](#) display.

URL	The IP address and port number for the server.
In Msgs/s	The amount of inbound messages for the topic, in number of messages per second.
In Total Msgs	The total number of inbound messages for the topic.
Out Msgs/s	The number of outbound messages per second.
Out Total Msgs	The total number of outbound messages for the topic since the server was started.
Pending Msgs	The number of currently pending messages for the topic.
Pending Msg Size	The amount of space, in bytes, pending messages use for the topic.
Active Durables	The number of currently active durables.
Consumers	The current number of consumers.
Durables	The number of active and inactive durables.
Fail Safe	When checked, the message is marked as failsafe delivery.
Flow Control Max Bytes	The maximum number of bytes allocated for use by flow control.
Global	When checked, the message is global and is routed to other servers.
In KB/s	The amount of inbound messages for the topic, in kilobytes per second.
In MB	The total number of inbound megabytes for the topic.
Max Bytes	The maximum size, in bytes, that the topic can store for delivery to each durable or non-durable online subscriber on the topic.
Max Msgs	The maximum number of messages allocated for use by the topic.
Out KB/s	The amount of outbound messages (in kilobytes) per second.
Out MB	The total amount of outbound messages for the topic, in megabytes, since the server was started.
Overflow Policy	Policy Indicates whether an overflow policy is set for the topic: 0 = No policy is set. 1 = A policy is set.
Secure	When checked, the topic is designated as secure and enforces permission policies.

Static	When checked, the topic has a static destination.
Subscribers	The number of subscribers for the topic.
Description	Descriptive text to help the administrator identify this resource.
Current In Total Msgs	Displays the change (delta) in inbound total messages from the previous cache refresh to the current cache refresh.
Current In Total Bytes	Displays the change (delta) in inbound total bytes from the previous cache refresh to the current cache refresh.
Current Out Total Msgs	Displays the change (delta) in outbound total messages from the previous cache refresh to the current cache refresh.
Current Out Total Bytes	Displays the change (delta) in outbound total bytes from the previous cache refresh to the current cache refresh.
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.

EMS Queues

These displays present several views of performance metrics for queues. Clicking **EMS Queues** from the left/navigation menu opens the “[TIBCO EMS Queues Table](#)” 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](#)”, 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](#)”, 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](#)”, 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](#)”, 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

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"** 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.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	
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

Filters




- EMS Server** The EMS Server selected from this drop-down menu populates all associated Queue data in this display.
- Filter by Queue Name** 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.
- Shown** The total number of currently active topics on the selected server, which is filtered by the **Data Collection > Metric Filters > Queues** field in the RTView Configuration Application.
- The default value for the **Queues** property is:
- ```
^(?!^\\$sys\\.|^\\$TMP\\.|^AMX_MGMT\\.|^EMSGMS\\.|^AMX_SV\\.|^_HAWK\\.|^_LOCAL\\.|^_HAWK\\.|^TMP\\.EMS)
```
- You can modify the filter value by editing the **Queues** property, which will override the default value. See ["Configuring Data Collection"](#) for more information.



**Total** The total number of queues on the selected server.

### Fields and Data

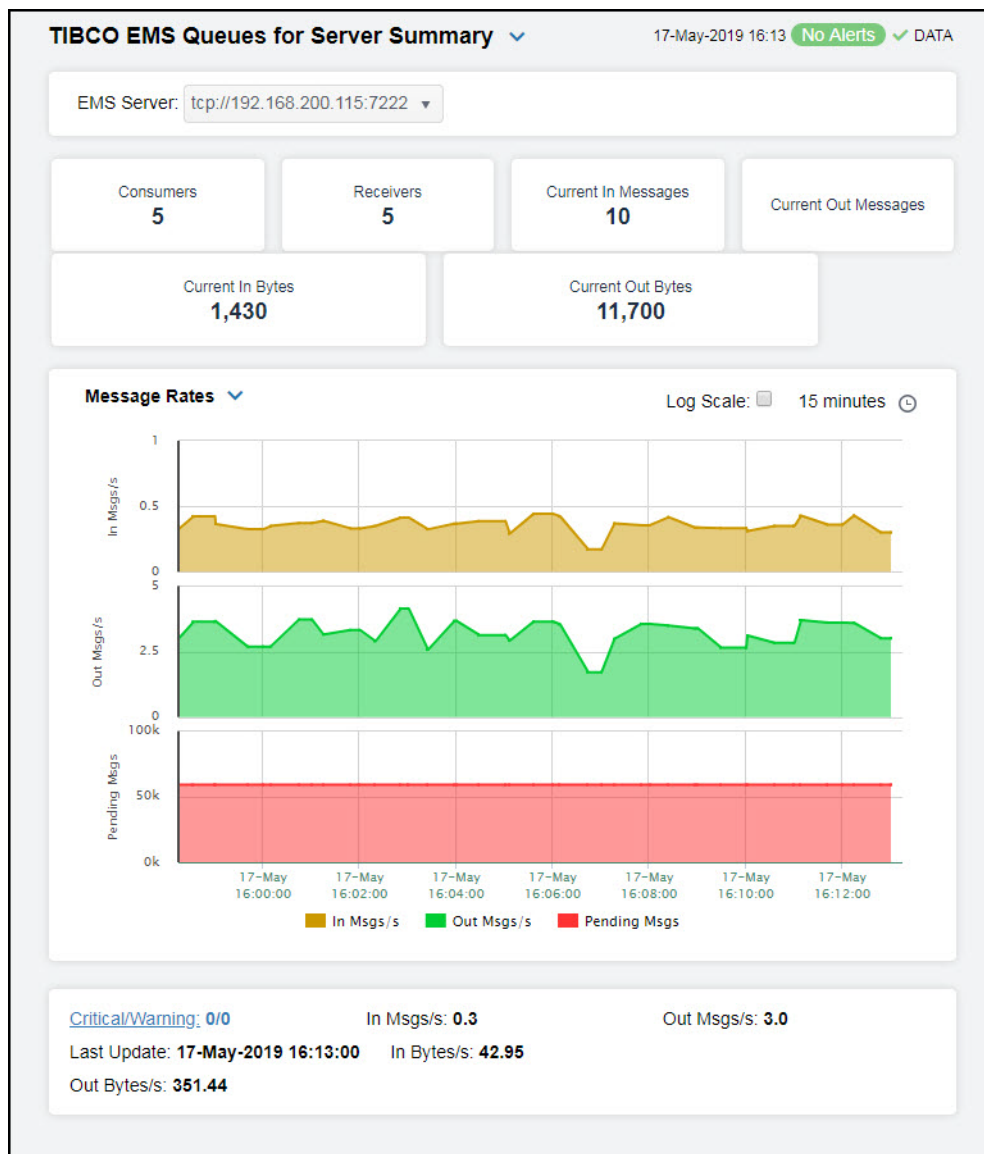
This display includes:

|                               |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Table</b>                  | This table describes all queues on the selected server. Double-click a row to view metrics for a single queue in the <a href="#">"TIBCO EMS Queue Summary"</a> display.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Queue Name</b>             | The name of the queue.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>URL</b>                    | The IP address and port number for the server.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Alert Level</b>            | The current alert level.<br> -- 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.<br> -- 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.<br> -- No alerts have exceeded an alert threshold, which have an Alert Severity value of <b>0</b> , and are shown in green. |
| <b>Alert Count</b>            | The number of current alerts.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>In Msgs/s</b>              | The number of inbound messages for the queue, per second.<br><b>Note:</b> This metric comes directly from the <b>tibjms.admin.DestinationInfo</b> class from TIBCO.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>In Total Msgs</b>          | The total number of inbound messages for the queue.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Out Msgs/s</b>             | The number of outbound messages for the queue, per second.<br><b>Note:</b> This metric comes directly from the <b>tibjms.admin.DestinationInfo</b> class from TIBCO.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Out Total Msgs</b>         | The total number of outbound messages for the queue.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Pending Msgs</b>           | The number of currently pending messages for the queue.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Consumers</b>              | The number of active and inactive consumers.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Fail Safe</b>              | When checked, the message is marked as failsafe delivery.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Flow Control Max Bytes</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 KB/s</b>                | The amount of inbound messages for the queue, in kilobytes per second.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>In MB</b>                  | The total amount of inbound messages for the queue, in megabytes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Max Bytes</b>              | The maximum amount of bytes allocated for use by the queue.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

|                             |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Max Msgs</b>             | The maximum number of messages allocated for use by the queue.                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Out KB/s</b>             | The amount of outbound messages for the queue, in kilobytes per second.                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Out MB</b>               | The total amount of outbound messages for the queue, in megabytes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Overflow Policy</b>      | Indicates whether an overflow policy is set for the queue:<br><b>0</b> = No policy is set.<br><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>Description</b>          | Descriptive text to help the administrator identify this resource.                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Pending Message Size</b> | The amount of space, in bytes, used by pending messages for the queue.                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Exclusive</b>            | When checked, the server sends all messages on this queue to one consumer.                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Max Redelivery</b>       | The maximum number of attempts for attempting redelivery of a message.                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Filter In Pattern</b>    | The string used to filter the data in the row.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Receivers</b>            | The number of receivers that receive queue message.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>In Msgs/s (TIBCO)</b>    | The number of inbound TIBCO messages for the queue, per second.<br>This metric comes directly from the <b>tibjms.admin.DestinationInfo</b> class from TIBCO.                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Out Msgs/s (TIBCO)</b>   | The number of outbound TIBCO messages for the queue, per second.<br>This metric comes directly from the <b>tibjms.admin.DestinationInfo</b> class from TIBCO.                                                                                                                                                                                                                                                                                                                                                                            |
| <b>In KB/s (TIBCO)</b>      | The amount of inbound TIBCO messages for the queue, in kilobytes per second.                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Out KB/s (TIBCO)</b>     | The amount of outbound TIBCO messages for the queue, in kilobytes per second.                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <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.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

## TIBCO EMS Queues for Server Summary

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"** 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.



**Filter**

**EMS Server** The EMS Server selected from this drop-down menu populates all associated queue data in this display.

**Fields and Data**

This display includes:

**Name** The name of the server selected in the **EMS Server** drop down list.

**Consumers** The number of consumers across all queues on the selected server.

**Receivers** The total number of receivers across all queues on the selected server.

**Current In Messages** The total number of inbound messages in the last update period across all queues on the selected server.

**Current Out Messages** The total number of outbound messages in the last update period across all queues on the selected server.

**Current In KB** The total number of inbound kilobytes in the last update period across all queues on the selected server.

**Current Out KB** The total number of outbound kilobytes in the last update period across all queues on the selected server.

**Trend Graphs****Message Rates**

Shows metrics for all queues on the selected server.

**In Msgs / sec** -- Traces the number of inbound messages for all queues, per second.

**Out Msgs / sec** -- Traces the number of outbound messages for all queues, per second.

**Pending Msgs** -- Traces the number of messages currently waiting to be processed.

**KB Rates**

Shows metrics for all queues on the selected server.

**In KB / sec** -- Traces the number of inbound messages for all queues, in kilobytes per second.

**Out KB / sec** -- Traces the number of outbound messages for all queues, in kilobytes per second.

**Pending Msg Size KB**-- Traces the amount of messages, in kilobytes, currently waiting to be processed.

**Pending Msgs**

Shows metrics for all queues on the selected server.

**Pending Msgs** -- Traces the number of messages currently waiting to be processed.


**Pending Msg Size KB**-- Traces the amount of messages, in kilobytes, currently waiting to be processed.

**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.

**Time Settings**

Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

**Critical/Warning**

The number of critical and warning alerts across all queues on the server.

**In Msgs/s**

The number of inbound messages for all queues, per second.

**Out Msgs/s**

The number of outbound messages for all queues, per second.

**Last Update**

The date and time of the last data update.

**In KB/s**

The number of inbound messages for all queues, in kilobytes per second.

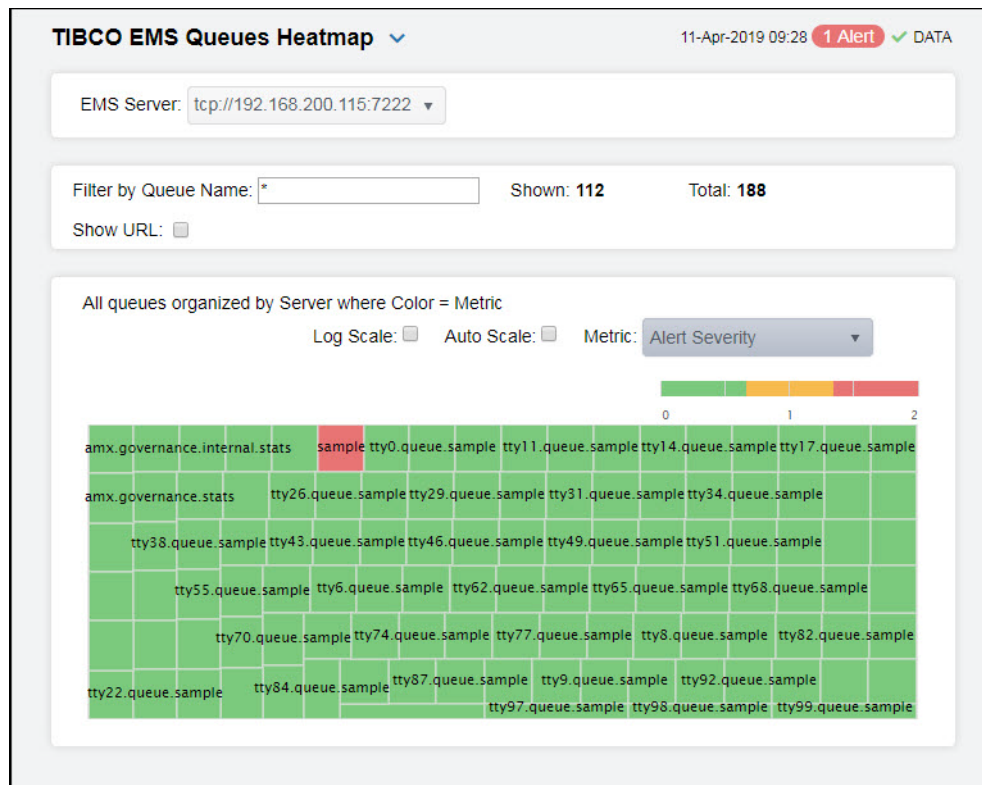
**Out KB/s**

The number of outbound messages for all queues, in kilobytes per second.

## TIBCO EMS Queues Heatmap



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” 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**.




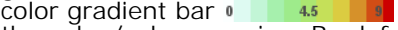

The heatmap is organized so that each rectangle represents a queue on the selected server. The rectangle color indicates the most critical alert state. Click on a node to drill-down to the “TIBCO EMS Queue Summary” display and view metrics for a particular queue. Toggle between the commonly accessed **Table** (link to the “TIBCO EMS Queues Table” display) and **Heatmap** displays. Mouse-over rectangles to view more details about the performance and status of each queue.



**EMS Server** The EMS Server selected from this drop-down menu populates all associated Queue data in this display.

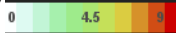
**Filter by Queue Name** Enter a string to show only queues with names that contain the string. For example, if you enter the string Madrid, all queues with Madrid in the queue name are shown in the table. If no entry is made, all queue names are shown. For most use cases, you can enter a portion of the queue name.

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Shown</b>      | <p>The total number of currently active topics on the selected server, which is filtered by the <b>Data Collection &gt; Metric Filters &gt; Queues</b> field in the RTView Configuration Application.</p> <p>The default value for the <b>Queues</b> property is:</p> <pre>^(?!^\\\\$sys\\. ^\\\\$TMP\\\\$\\. ^AMX_MGMT\\. ^EMSGMS\\. ^AMX_SV\\. ^_HAWK\\. ^_LOCAL\\._HAWK\\. ^TMP\\.EMS)</pre> <p>You can modify the filter value by editing the <b>Queues</b> property, which will override the default value. See <a href="#">"Configuring Data Collection"</a> for more information.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Total</b>      | The total number of queues on the selected 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.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Auto Scale</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>     | <p>Select the metric driving the heatmap display. The default is <b>Alert Severity</b>. 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 <a href="#">"TIBCO EMS Queue Summary"</a> display for a detailed view of metrics for that particular queue.</p> <p><b>Alert Severity</b>      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> <ul style="list-style-type: none"> <li>-- 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.</li> <li><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.</li> <li>-- 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.</li> </ul> <p><b>Alert Count</b>      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 0 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>                                                                                                                                                                                                                                                                                                                                                                                              |
| <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 0 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>                                                                                                                                                                                                                                                                                                                                                                                              |
| <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 0 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 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 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 0 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 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 0 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 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 Scale** option does not impact this metric.

## TIBCO EMS Queue Summary

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"](#) 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.



## Filters

- 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.
- Queue** Select a queue from the drop-down menu. The selection made here populates this display.


## Fields and Data

This display includes:

- Pending Messages** The total number of messages for the selected queue currently waiting to be processed.
- Consumers** The number of consumers currently interacting with the queue.
- In Msgs/s** The number of inbound messages, per second, for the selected queue.
- Out Msgs/s** The number of outbound messages, per second, for the selected queue.
- In KB/s** The size of the inbound messages, in kilobytes per second, for the selected queue.
- Out KB/s** The size of outbound messages, in kilobytes per second, for the selected queue.
- Trend Graphs**
- Message Rates**  
Shows metrics for the selected queue on the specified server.
- In Msgs / sec** -- Traces the number of inbound messages, per second.
  - Out Msgs / sec** -- Traces the number of outbound messages, per second.
  - Pending Msgs** -- Traces the number of messages currently waiting to be processed.
- Message Flows**  
Shows metrics for the selected queue on the specified server.
- In Msgs**-- Traces the number of inbound messages.
  - Out Msgs** -- Traces the number of outbound messages.
  - Pending Msgs** -- Traces the number of messages currently waiting to be processed.
- Use Rates** When this check box is selected, the inbound and outbound message rates (**In Msgs/sec** and **Out Msgs/sec**) display in the trend graph. When this check box is not selected, the delta inbound and outbound messages (**Delta In Msgs** and **Delta Out Msgs**) display in the trend graph.
- 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.

**Time Settings**

Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

|                                |                                                                                               |
|--------------------------------|-----------------------------------------------------------------------------------------------|
| <b>Receivers</b>               | The number of consumers currently receiving messages from the queue.                          |
| <b>Pending Message Size MB</b> | The size, in megabytes, of messages for the selected queue currently waiting to be processed. |
| <b>Global</b>                  | When checked, the message is global and is routed to other servers.                           |
| <b>Secure</b>                  | When checked, the queue is designated as secure and enforces permission policies.             |
| <b>Critical/Warning</b>        | The number of critical and warning alerts on the queue.                                       |
| <b>Max MB</b>                  | The maximum amount of memory, in megabytes, allocated for use by the queue.                   |
| <b>Static</b>                  | When checked, the queue has a static destination.                                             |
| <b>Exclusive</b>               | When checked, the server sends all messages on this queue to one consumer.                    |
| <b>Max Msgs</b>                | The maximum number of messages allocated for the queue.                                       |
| <b>FC Max MB</b>               | The maximum amount of memory, in megabytes, allocated for flow control use by the queue.      |
| <b>Fail Safe</b>               | When checked, the message is marked as failsafe delivery.                                     |

|                        |                                                                                                             |
|------------------------|-------------------------------------------------------------------------------------------------------------|
| <b>Overflow Policy</b> | Indicates whether an overflow policy is set for the queue:<br>0 = No policy is set.<br>1 = A policy is set. |
| <b>Last Update</b>     | The date and time of the last data update.                                                                  |

## TIBCO EMS Queue Detail By Server

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"](#) display, where you can view additional data for that particular queue on that particular server.

| 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              |

### Filter

**Queue** The Queue selected from this drop-down menu populates this display.

### Fields and Data

This display includes:

|                               |                                                                                                                                                                              |
|-------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Count</b>                  | The number of rows found based on the filter and displayed in the table.                                                                                                     |
| <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="#">“TIBCO EMS Queue Summary”</a> display. |
| <b>URL</b>                    | The URL of the server.                                                                                                                                                       |
| <b>In Msgs/s</b>              | The amount of inbound messages for the queue, in number of messages per second.                                                                                              |
| <b>In Total Msgs</b>          | The total number of inbound messages for the queue.                                                                                                                          |
| <b>Out Msgs/s</b>             | The number of outbound messages per second.                                                                                                                                  |
| <b>Out Total Msgs</b>         | The total number of outbound messages since the server was started.                                                                                                          |
| <b>Pending Msgs</b>           | The number of currently pending messages.                                                                                                                                    |
| <b>Consumers</b>              | The number of active and inactive consumers.                                                                                                                                 |
| <b>Fail Safe</b>              | When true, the message is marked as failsafe delivery.                                                                                                                       |
| <b>Flow Control Max Bytes</b> | The maximum number of bytes allocated for use by flow control.                                                                                                               |
| <b>Global</b>                 | When true, the message is global and is routed to other servers.                                                                                                             |
| <b>In KB/s</b>                | The amount of inbound messages for the queue, in kilobytes per second.                                                                                                       |
| <b>In MB</b>                  | The total number of inbound megabytes for the queue.                                                                                                                         |
| <b>Max Bytes</b>              | The maximum amount of bytes allocated for use by the queue.                                                                                                                  |
| <b>Max Msgs</b>               | The maximum number of messages allocated for use by the queue.                                                                                                               |
| <b>Out KB/s</b>               | The amount of outbound messages (in kilobytes) per second.                                                                                                                   |
| <b>Out MB</b>                 | The total amount of outbound messages, in megabytes, since the server was started.                                                                                           |
| <b>Overflow Policy</b>        | Indicates whether an overflow policy is set for the queue:<br><b>0</b> = No policy is set.<br><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>Description</b>            | Descriptive text to help the administrator identify this resource.                                                                                                           |
| <b>Current In Total Msgs</b>  | The total number of inbound messages in the last update period for the queue.                                                                                                |

|                                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Current In Total Bytes</b>  | The total number of inbound bytes in the last update period for the queue.                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Current Out Total Msgs</b>  | The total number of outbound messages in the last update period for the queue.                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Current Out Total Bytes</b> | The total number of outbound bytes in the last update period for the queue.                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Pending Msg Size</b>        | The amount of space, in bytes, pending messages use for the queue.                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Exclusive</b>               | When checked, the server sends all messages on this queue to one consumer.                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Max Redelivery</b>          | The maximum number of attempts for attempting redelivery of a message.                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Filter In Pattern</b>       | The string used to filter the data in the row.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Receivers</b>               | The number of receivers of queue messages.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <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 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.                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

## EMS Routes

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](#)" 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](#)" display, which shows metrics and trend data for a particular route on a particular EMS Server.

## TIBCO EMS Routes

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](#)" 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://192.168.200.171:601 | EMS-SLDEMOS2-6010 | 0         |         |

Last Update: **11-Apr-2019 11:19:53**

### Filter

**EMS Server** The EMS Server selected from this drop-down menu populates all associated Routes data in this display.

### Fields and Data

This display includes:

**Status** The current status of the server.

**Totals For Server** Shows metrics for all server routes on the selected server.

**In Msgs / sec** The number of inbound messages, per second.

**In Total Msgs** The total number of inbound messages.

**Out Msgs / sec** The number of outbound messages, per second.

**Out Total Msgs** 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.

**Remote URL** The URL of the remote server.

**Remote Name** The name of the remote server.

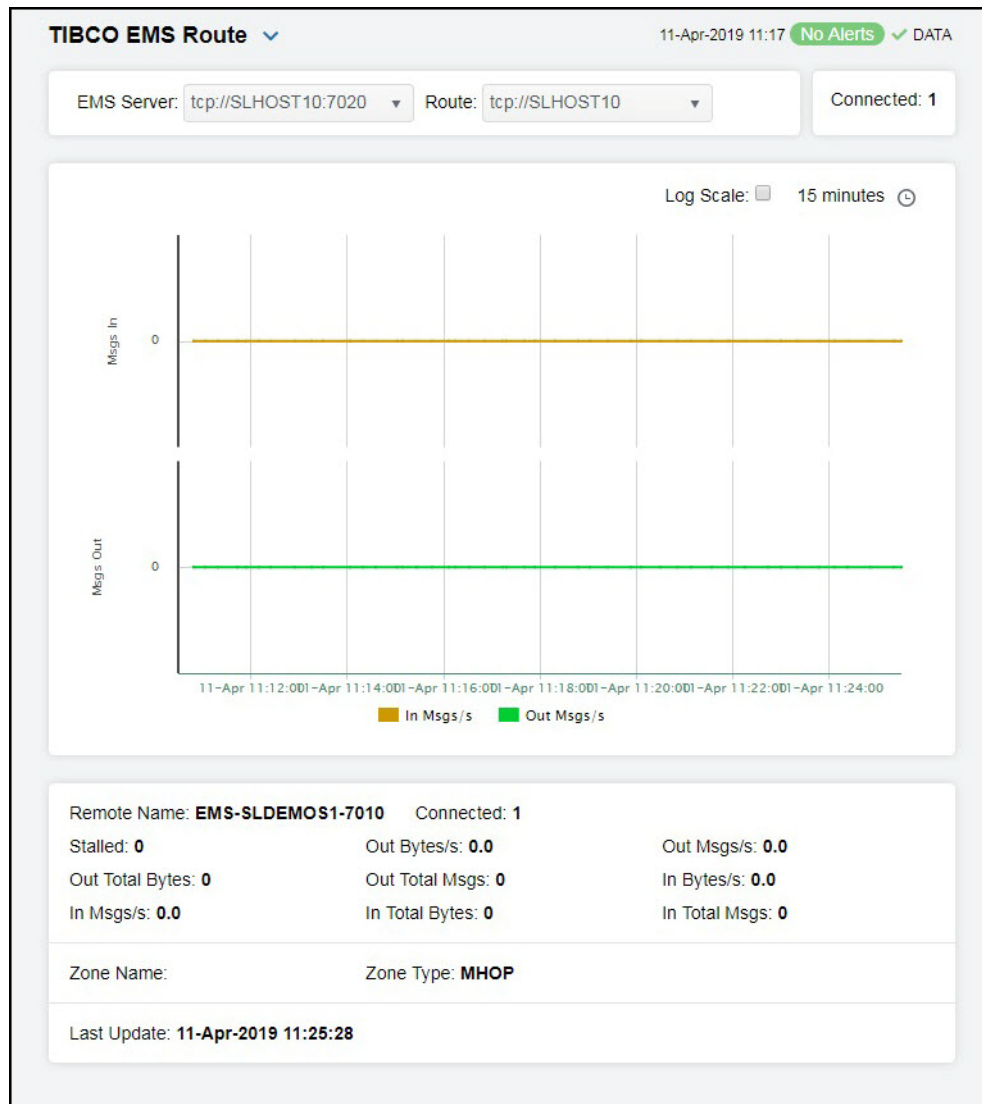
**Connected** When checked, the server route is connected.



|                          |                                                                                                                                    |
|--------------------------|------------------------------------------------------------------------------------------------------------------------------------|
| <b>Stalled</b>           | Indicates whether the IO flow stalled on the route.<br>A value of <b>0</b> (zero) = not stalled.<br>A value of <b>1</b> = stalled. |
| <b>In Bytes/s</b>        | The rate of inbound data in bytes, per second.                                                                                     |
| <b>In Msgs/s</b>         | The rate of inbound messages in number of messages per second.                                                                     |
| <b>In Total Bytes</b>    | The total number of inbound bytes.                                                                                                 |
| <b>In Total Msgs</b>     | The total number of inbound messages.                                                                                              |
| <b>Out Bytes/s</b>       | The rate of outbound data in bytes per second.                                                                                     |
| <b>Out Msgs/s</b>        | The rate of outbound messages in number of messages per second.                                                                    |
| <b>Out Total Bytes</b>   | The total number of outbound bytes.                                                                                                |
| <b>Out Total Msgs</b>    | The total number of outbound messages.                                                                                             |
| <b>Zone Name</b>         | The name of the zone for the route.                                                                                                |
| <b>Zone Type</b>         | Indicates a multi-hop or one-hop zone.                                                                                             |
| <b>Active</b>            | Indicates whether the server route is currently transferring data:<br><b>1</b> = true<br><b>0</b> = false                          |
| <b>Inactive</b>          | Indicates whether the server route is currently transferring data:<br><b>1</b> = true<br><b>0</b> = false                          |
| <b>Suspended</b>         | Indicates whether outbound messages to the route have been suspended:<br><b>1</b> = true<br><b>0</b> = false                       |
| <b>remoteURLNoCommas</b> | The IP address and name for the remote connection.                                                                                 |
| <b>Time Stamp</b>        | The date and time of the last data update.                                                                                         |
| <b>Last Update</b>       | The date and time of the last data update.                                                                                         |

## TIBCO EMS Route

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.



### Filter

**EMS Server** The EMS Server selected from this drop-down menu populates all associated Routes data in this display.

**Route** Select the route for which you want to view data in the display.

### Fields and Data

This display includes:

**Connected** The number of routes connected.

**Trend Graphs** Shows message data for the selected route.

**In Msgs/s** -- Traces the number of inbound messages, per second.


**Out Msgs/s** -- Traces the number of outbound messages, 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.

**Time Settings**

Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

|                        |                                                                                                                                    |
|------------------------|------------------------------------------------------------------------------------------------------------------------------------|
| <b>Remote Name</b>     | The name of the remote server.                                                                                                     |
| <b>Stalled</b>         | Indicates whether the IO flow stalled on the route.<br>A value of <b>0</b> (zero) = not stalled.<br>A value of <b>1</b> = stalled. |
| <b>Out Total Bytes</b> | The total number of outbound bytes.                                                                                                |
| <b>In Msgs/s</b>       | The rate of inbound messages in number of messages per second.                                                                     |
| <b>Connected</b>       | The number of routes connected.                                                                                                    |
| <b>Out Bytes/s</b>     | The rate of outbound data in bytes per second.                                                                                     |

|                       |                                                                 |
|-----------------------|-----------------------------------------------------------------|
| <b>Out Total Msgs</b> | The total number of outbound messages.                          |
| <b>In Total Bytes</b> | The total number of inbound bytes.                              |
| <b>Out Msgs/s</b>     | The rate of outbound messages in number of messages per second. |
| <b>In Bytes/s</b>     | The rate of inbound data in bytes, per second.                  |
| <b>In Total Msgs</b>  | The total number of inbound messages.                           |
| <b>Zone Name</b>      | The name of the zone for the route.                             |
| <b>Zone Type</b>      | Indicates a multi-hop or one-hop zone.                          |
| <b>Last Update</b>    | The date and time of the last data update.                      |

## EMS Producers

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](#)” 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](#)” display, which shows metrics and trend data for a particular producer on a particular EMS Server.

## TIBCO EMS Producers

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](#)” display. Clicking the **Queues** link in the bottom portion of the display opens the “[TIBCO EMS Queues Table](#)” 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**                      Queues: **242**                      Durables: **5**  
 Producers: **30**                      Consumers: **88**                      Routes: **1**

Last Update: **11-Apr-2019 13:38:15**

## Filters

- EMS Server**                      The EMS Server selected from this drop-down list displays a list of the currently connected Producers. The field to the right of the EMS Server drop down displays the status of the server.
- Client ID**                      Filter field that allows you to filter the list of producers by client ID.
- RegEx**                      Select this toggle to use a regular expression for the **Client ID** filter field.
- Destination**                      Filter field that allows you to filter the list of producers by destination name.
- RegEx**                      Select this toggle to use a regular expression for the **Destination** filter field.

## Fields and Data

This display includes:

- Count**                      The number of currently connected producers on the server.
- Msgs/s**                      The number of messages, per second, for the producer.
- Msgs Total**                      The total number of messages for the producer.

|                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Bytes/s</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>Producers for Server 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="#">"TIBCO EMS Producer"</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 Name</b>           | The name of the destination.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Msgs/s</b>                     | The number of messages, per second, for the producer.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Msgs Total</b>                 | The total number of messages for the producer.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Bytes/s</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 Name</b>                  | The user name.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Host</b>                       | The name of the host.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <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>Create Time</b>                | The amount of time, in milliseconds, since the producer was created.                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Destination Type</b>           | The configured destination type.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <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>Topics</b>                     | The total number of topics on the server (pulled directly from the TIBCO API).                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Producers</b>                  | The total number of producers (pulled directly from the TIBCO API).                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Queues</b>                     | The total number of queues on the server (pulled directly from the TIBCO API).                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Consumers</b>                  | The total number of consumers (pulled directly from the TIBCO API).                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Durables</b>                   | The total number of durables on the server (pulled directly from the TIBCO API).                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Routes</b>                     | The total number of routes on the server (pulled directly from the TIBCO API).                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Last Update</b>                | The date and time of the last data update (pulled directly from the TIBCO API).                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

## TIBCO EMS Producer

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.



### Filters


**EMS Server** The selected EMS Server populates the Producer ID/ Client ID drop-down menu with associated Producer IDs/Client IDs.

**ID** Drop-down menu containing the Producer IDs/Client IDs.

### Fields and Data

This display includes:

- Trend Graph** Shows message data for the selected producer.
- Bytes/s** -- Traces the size of messages for the producer, in bytes.
  - Msgs/s** -- Traces the number of messages for the producer, 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.
- Time Settings** Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

- Destination Type** The configured destination type.
- Destination Name** The name of the destination.
- Client ID** A unique string identifier assigned to each client.
- Conn ID** A unique string identifier assigned to each connection.
- Msgs/s** The number of messages, per second, for the producer.



|                    |                                                              |
|--------------------|--------------------------------------------------------------|
| <b>Total Bytes</b> | The total size of messages, in bytes, for the producer.      |
| <b>User Name</b>   | The user name.                                               |
| <b>Create Time</b> | The date and time when the producer was created.             |
| <b>Msgs Total</b>  | The total number of messages for the producer.               |
| <b>Session ID</b>  | A unique string identifier assigned to each session.         |
| <b>Bytes/s</b>     | The size of messages, in bytes per second, for the producer. |
| <b>Last Update</b> | The date and time of the last data update.                   |

## EMS Consumers

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"](#) 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"](#) display, which shows metrics and trend data for a particular consumer on a particular EMS Server.

## TIBCO EMS Consumers

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"](#) display. Clicking the **Queues** link in the bottom portion of the display opens the ["TIBCO EMS Queues Table"](#) display.

**TIBCO EMS Consumers** 11-Apr-2019 13:55 No Alerts DATA

EMS Server:  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 | Consumer 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,              |
| 193104 |           | c_l_payload_queue    | 0.0             |                     |
| 193120 |           | c_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**

## Filters

- EMS Server** The EMS Server selected from this drop-down list displays a list of the currently connected Consumers. The field to the right of the EMS Server drop down displays the status of the server.
- Client ID** Filter field that allows you to filter the list of consumers by client ID.
- RegEx** Select this toggle to use a regular expression for the **Client ID** filter field.
- Destination** Filter field that allows you to filter the list of consumers by destination name.
- RegEx** Select this toggle to use a regular expression for the **Destination** filter field.

## Fields and Data

This display includes:

- Count** The number of currently connected producers on the server.
- Consumer Msgs/s** The number of messages, per second, for the consumer.
- Consumer Total Msgs** The total number of messages for the consumer.

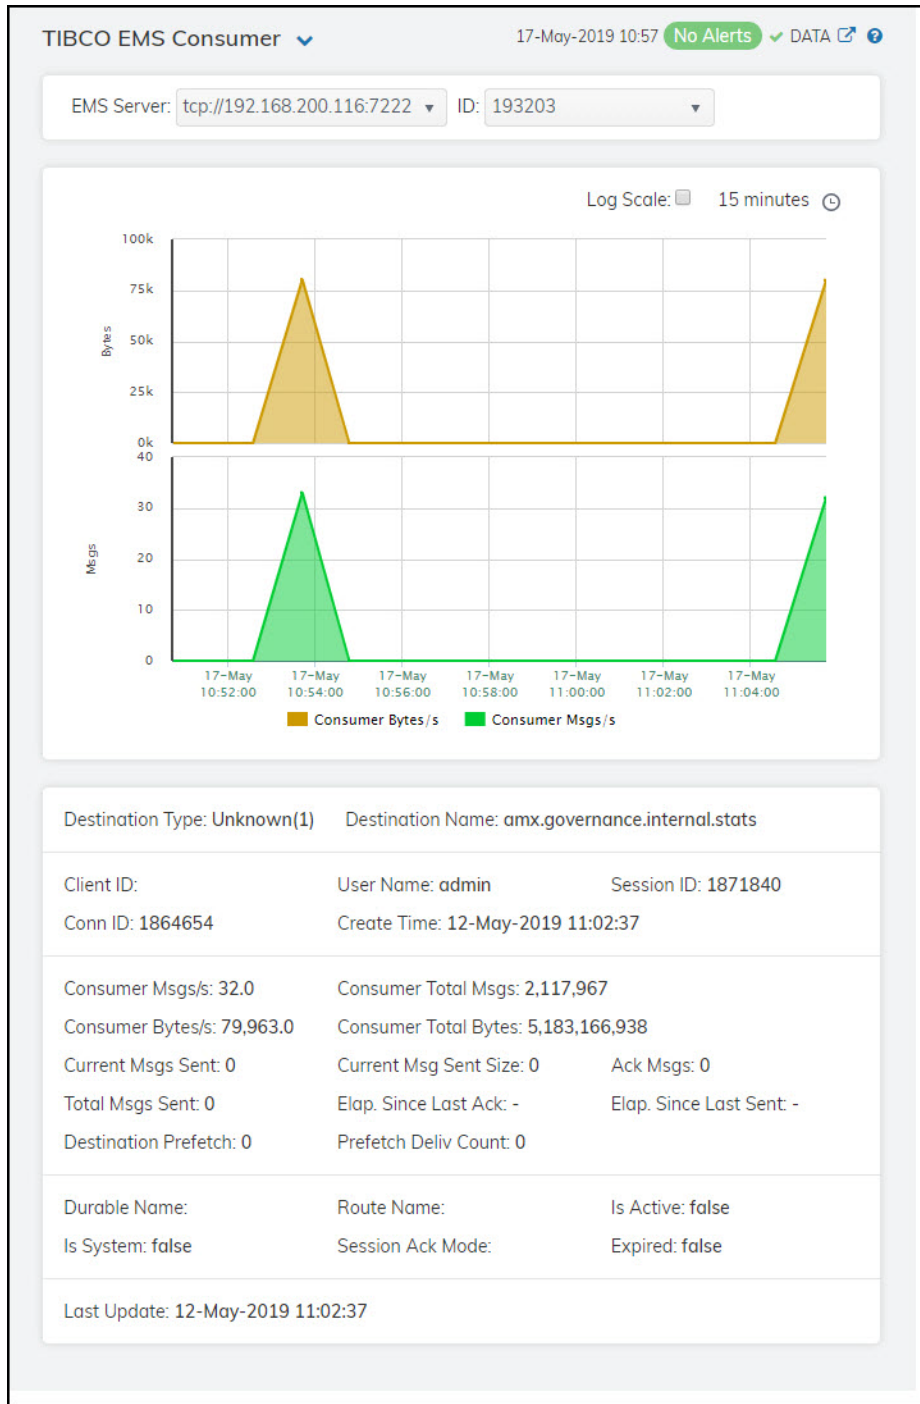
|                                   |                                                                                                                                                                                                                                                                                                                |
|-----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Consumer Bytes/s</b>           | The amount of messages, in bytes per second, for the consumer.                                                                                                                                                                                                                                                 |
| <b>Consumer Total Bytes</b>       | The total size of messages, in bytes, for the consumer.                                                                                                                                                                                                                                                        |
| <b>Consumers for Server 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 "TIBCO EMS Consumer" 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>Destination Name</b>           | The name of the destination.                                                                                                                                                                                                                                                                                   |
| <b>Consumer Msgs/s</b>            | The number of messages, per second, for the consumer.                                                                                                                                                                                                                                                          |
| <b>Consumer Total Msgs</b>        | The total number of messages for the consumer.                                                                                                                                                                                                                                                                 |
| <b>Consumer Bytes/s</b>           | The size of messages, in bytes per second, for the consumer.                                                                                                                                                                                                                                                   |
| <b>Consumer Total Bytes</b>       | The total size of messages, in bytes, for the consumer.                                                                                                                                                                                                                                                        |
| <b>User Name</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 Msgs Sent</b>             | The number of messages sent to the consumer that were not yet acknowledged by the consumer's session.<br><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>Curr Msg Sent Size</b>         | The combined size of messages sent to the consumer that were not yet acknowledged by the consumer's session.<br><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>Ack Msgs</b>                   | The total number of messages that have been sent to the consumer and have been acknowledged by the consumer's session.<br><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 Msgs Sent</b>            | The total number of messages sent to the consumer since the consumer was created.<br><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>Elap. 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.<br><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>Elap. 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.<br><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>Destination Prefetch</b>  | The actual destination prefetch value used by the server at runtime.<br><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>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.<br><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>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.<br><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>Is Active</b>             | If true, the consumer is active and can receive messages from the server.<br><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>Is System</b>             | If true, the consumer was automatically created by the system.<br><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>Session Ack Mode</b>      | Lists the consumer's session acknowledge mode as a constant defined in <b>TibjmsAdmin</b> .<br><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>Create Time</b>           | The amount of time, in milliseconds, since the consumer was created.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Destination Type</b>      | The type of 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 > (Project Name) > <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>Topics</b>                | The total number of topics on the server (pulled directly from the TIBCO API).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Producers</b>             | The total number of producers (pulled directly from the TIBCO API).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Queues</b>                | The total number of queues on the server (pulled directly from the TIBCO API).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

- Consumers** The total number of consumers (pulled directly from the TIBCO API).
- Durables** The total number of durables on the server (pulled directly from the TIBCO API).
- Routes** The total number of routes on the server (pulled directly from the TIBCO API).
- Last Update** The date and time of the last data update (pulled directly from the TIBCO API).

## TIBCO EMS Consumer

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.



**Filters**


**EMS Server**      The selected EMS Server populates the Consumer ID/ Client ID drop-down menu with associated Producer IDs/Client IDs.

**ID**                Drop-down menu containing the Consumer IDs/Client IDs.

## Fields and Data

This display includes:

- 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.
- Time Settings** Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

- Destination Type** The configured destination type.
- Destination Name** The name of the destination.
- Client ID** A unique string identifier assigned to each client.
- Conn ID** A unique string identifier assigned to each connection.

|                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>User Name</b>             | The user name.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Create Time</b>           | The amount of time, in milliseconds, since the consumer was created.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Session ID</b>            | A unique string identifier assigned to each session.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Consumer Msgs/s</b>       | The number of messages, per second, for the consumer.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Consumer Total Bytes</b>  | The total size of messages, in bytes, for the consumer.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Ack Msgs</b>              | The total number of messages that have been sent to the consumer and have been acknowledged by the consumer's session.<br>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>Elap. 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.<br>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>Consumer Total Msgs</b>   | The total number of messages for the consumer.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Current Msgs Sent</b>     | The number of messages sent to the consumer that were not yet acknowledged by the consumer's session.<br>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 Msgs Sent</b>       | The total number of messages sent for the consumer.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Destination Prefetch</b>  | The actual destination prefetch value used by the server at runtime.<br>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>Consumer Bytes/s</b>      | The size of messages, in bytes per second, for the consumer.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Current Msg Sent Size</b> | The combined size of messages sent to the consumer that were not yet acknowledged by the consumer's session.<br><b>Note: The sl.rtvview.jmsadm.queryCIDetails property must be set to true in your sample.properties file to see this column.</b>                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Elap. 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.<br>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>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.<br>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>Durable Name</b>          | The name of the durable.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Is System</b>             | If true, the consumer was automatically created by the system.<br>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>Route Name</b>       | The queue owner server name if the consumer if the consumer's destination is a routed queue.<br>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>Session Ack Mode</b> | Lists the consumer's session acknowledge mode as a constant defined in <b>TibjmsAdmin</b> .<br>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>Is Active</b>        | If true, the consumer is active and can receive messages from the server.<br>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>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>Last Update</b>      | The date and time of the last data update.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

## EMS Durables

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**" 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**" display, which shows metrics and trend data for a particular durable on a particular EMS Server.

## TIBCO EMS Durables

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**" display.

**TIBCO EMS Durables** 11-Apr-2019 14:19 No Alerts DATA

EMS Server:  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**

## Filter

**EMS Server**      The EMS Server selected from this drop-down menu populates all associated Durables data in this display. The field to the right of the EMS Server drop down displays the status of the server.

## Fields and Data

This display includes:

**Total Pending Msgs**      The total number of pending messages for the durable.

**Total Pending Size**      The total amount of pending messages, in bytes, for the selected durable.

**Durables for Server Table**      This table shows metrics for each durable on the selected server.

**Name**      The name of the durable.

**Topic**      The name of the topic.

**Active**      Indicates whether the durable is active.

**Client ID**      A unique string identifier assigned to each client.

|                         |                                                                                                                                                                                                                                                                       |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Consumer ID</b>      | A unique string identifier assigned to each consumer.                                                                                                                                                                                                                 |
| <b>No Local Enabled</b> | Indicates whether the subscriber receives messages from all connections or its local connection.<br><b>True</b> -- The subscriber does not receive messages sent from its local connection.<br><b>False</b> -- The subscriber receives messages from all connections. |
| <b>Pending Msgs</b>     | The total number of pending messages for the selected durable.                                                                                                                                                                                                        |
| <b>Pending Msg 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>User Name</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>Last Data Update</b> | The date and time of the last data update.                                                                                                                                                                                                                            |

## TIBCO EMS Durable

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.



### Filter


- EMS Server** The EMS Server selected from this drop-down menu populates all associated Durables data in this display. The field to the right of the EMS Server drop down displays the status of the server.
- Durable** Select the durable for which you want to view data.
- Client ID** Select the client ID for which you want to view data.

### Fields and Data

This display includes:

- Total Pending Msgs** The total number of pending messages for the durable.
- Total Pending Size** The total amount of pending messages, in bytes, for the selected durable.

- Trend Graph** Shows message data for the selected consumer.
- Pending Msgs** -- Traces the number of pending messages for the durable.
- 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.
- Time Settings** Select a time range from the drop down menu varying from **5 Minutes** to **Last 7 Days**. By default, the time range end point is the current time.



To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated calendar and clock icons. You can also enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2018 12:24 PM. Click the **now** toggle to reset the time range end point to the current time.

- Topic** The name of the topic.
- No Local Enabled** Indicates whether the subscriber receives messages from all connections or its local connection.  
**True** -- The subscriber does not receive messages sent from its local connection.  
**False** -- The subscriber receives messages from all connections.
- Selector** Indicates that the subscriber only receives messages that match this selector.
- Active** Indicates whether the durable is active.
- Pending Msgs** The total number of pending messages for the selected durable.

|                         |                                                                           |
|-------------------------|---------------------------------------------------------------------------|
| <b>User Name</b>        | The name of the user of this durable subscriber.                          |
| <b>Consumer ID</b>      | A unique string identifier assigned to each consumer.                     |
| <b>Pending Msg Size</b> | The total amount of pending messages, in bytes, for the selected durable. |
| <b>Last Update</b>      | The date and time of the last data update.                                |

---

## TIBCO FTL

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”](#) (shown below), and the following Views which can be found under **Components** tab > **Middleware**.

- [“FTL Servers”](#)
- [“FTL Clients”](#)

### TIBCO FTL Overview

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.

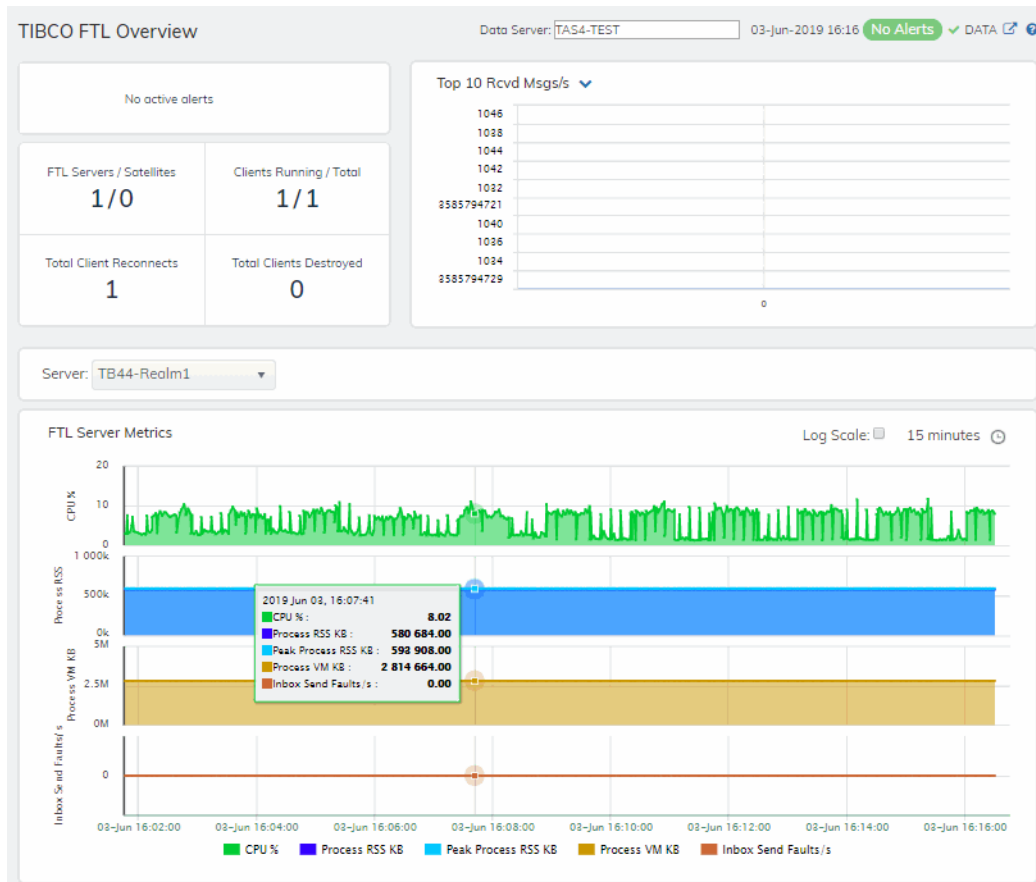
You can select the RTView Data Server for which you want to see data and easily view the current data for that Data Server including:

- 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”](#) 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


These displays present performance metrics and alert status for all FTL servers. Displays are:

- [“FTL Servers Heatmap”](#): Heatmap shows server and alert status for all FTL servers in all realms.
- [“FTL Servers Table”](#): Table shows all available utilization metrics for all FTL servers.
- [“FTL Server Summary”](#): Current and historical metrics for a single FTL server.

## FTL Servers Heatmap

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.




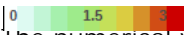
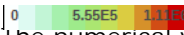
Each heatmap rectangle represents a server. The rectangle color indicates the most critical alert state. You can click on a rectangle to drill-down to the “[FTL Server Summary](#)” display and view metrics for that server. Clicking on the  icon next to the display title toggles between the commonly accessed **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.



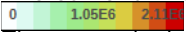


## Fields and Data

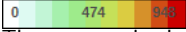
This display includes:

|                       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <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”</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

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”](#) 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

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”](#) display, where you can view additional server data.

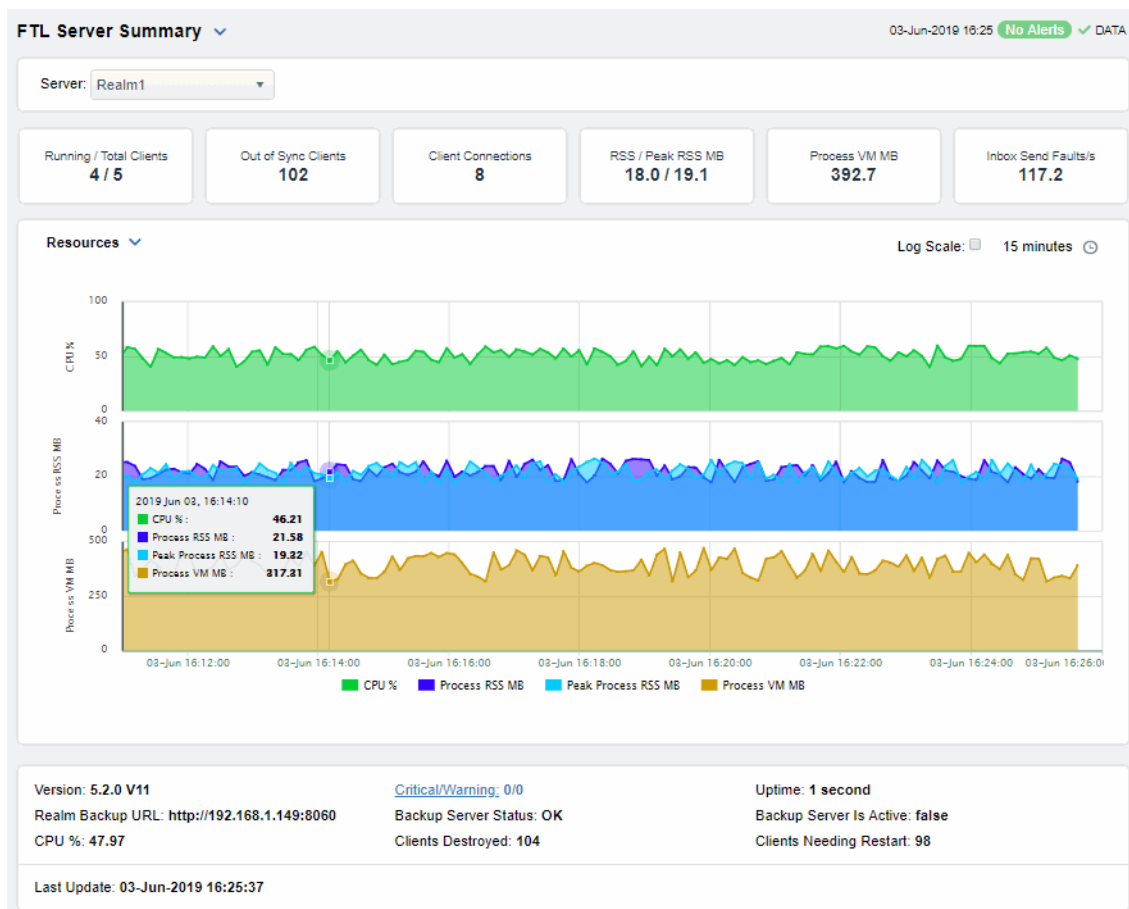
Choose from one of two trend graphs:

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.



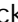
## FTL Clients

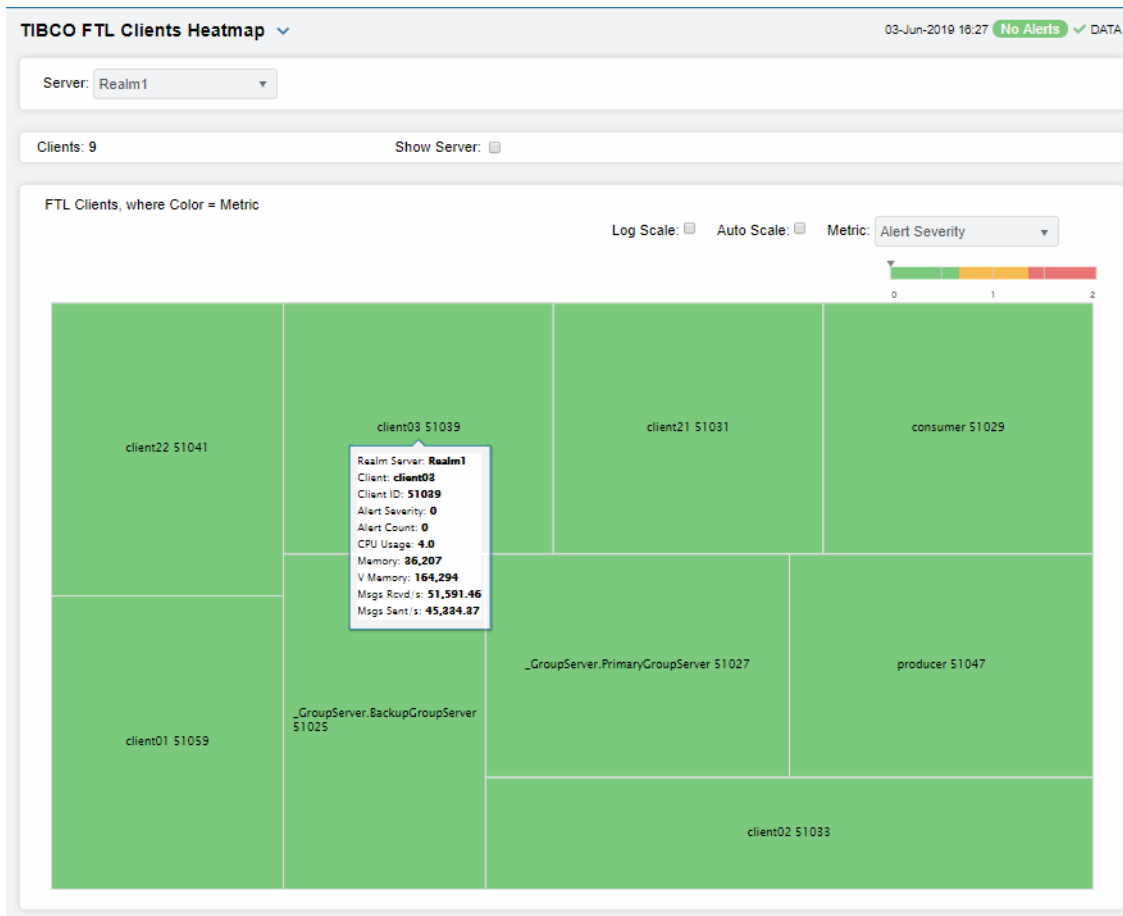
These displays present performance metrics and alert status for FTL clients. Available displays are:

- [“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.
- [“Single Client Summary”](#): Displays current and historical metrics for a single FTL client.

### All Clients Heatmap

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”](#) 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.




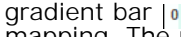
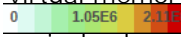
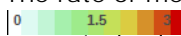

**Fields and Data**

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"</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>TftIClientCpuUsageHigh</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>TftIClientMemoryHigh</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 <u>virtual memory</u> 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>TftIClientVirtualMemoryHigh</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>TftIClientMsgsRcvdRateHigh</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>TftIClientMsgsSentRateHigh</b>. The middle value in the gradient bar indicates the middle value of the range.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

## All Clients Table

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”** 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”**, 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

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”](#) display, where you can view and compare data against other clients.

Select a Server and Client from the drop-down menus. There are two trend graphs you can choose from:

- 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.







## CHAPTER 9 RTView Manager

Use the RTView Manager application to track the health of your RTView Enterprise system: RTView and Tomcat processes, the historian and RTView data servers. RTView Manager runs as a process, separately from your RTView monitoring system, has its own URL / console, as well as its own data server, database, alert notification system and historian.

For details about configuring RTView Manager, refer to the *RTView Enterprise Configuration Guide*.

This chapter contains:

- [“Displays”](#): Describes [“Tomcat Displays”](#), [“JVM Processes Displays”](#), [“RTView Servers Displays”](#) and [“Drilldowns’ Displays”](#) which are accessed from the navigation tree (left panel).
- [“Alerts Displays”](#): Describes the display that is used to manage alerts.
- [“Admin Displays”](#): Describes displays that administrators use to manage alert thresholds.
- [“Alerts for RTView Manager”](#): Describes RTView Manager alerts.

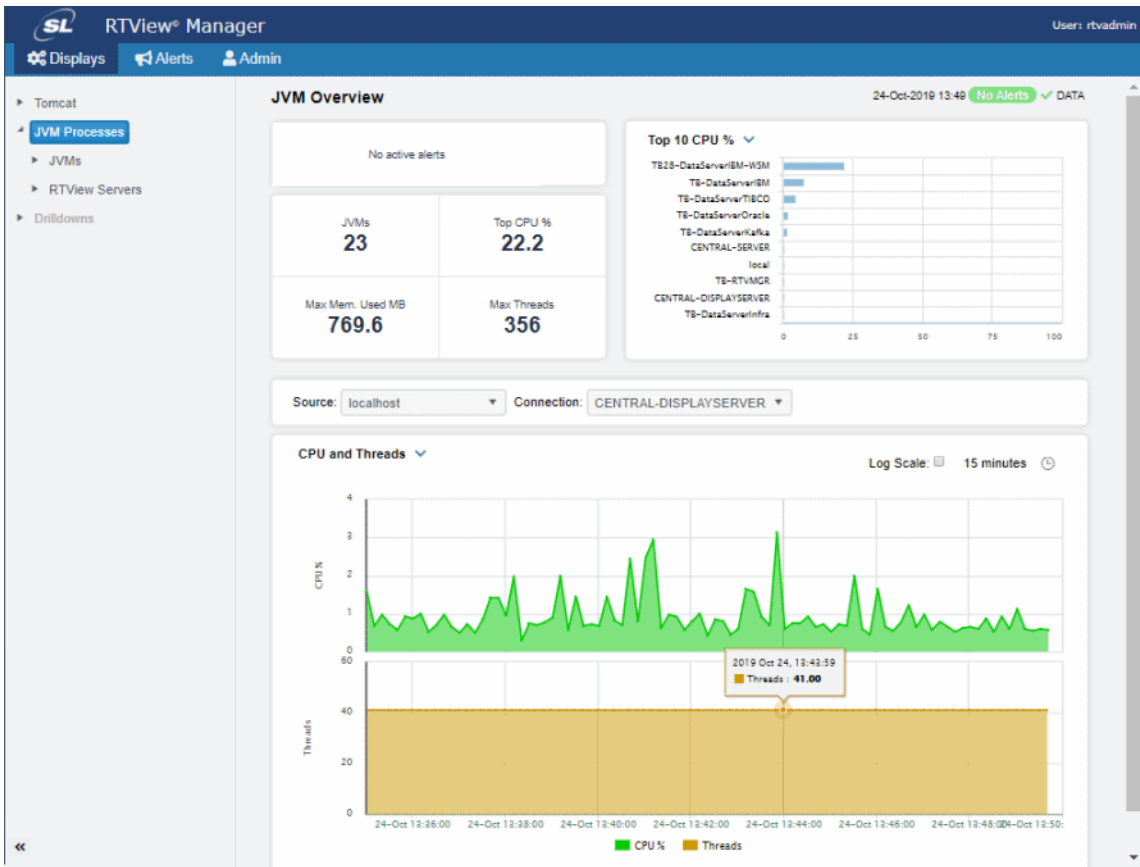
### Login to RTView Manager

To access RTView Manager, start RTViewCentral (if not currently running), then browse to one of the following URLs and login (username/password are rtvadmin/rtvadmin):

[http://<ip\\_address>:3070/rtview-manager](http://<ip_address>:3070/rtview-manager) if you are running on Jetty.

<http://localhost:8068/rtview-manager> if you are using Tomcat.

The RTView Manager main console opens.



---

## Displays

The displays that come with RTView Manager are organized by the following Views in the navigation tree:

- **“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.
- **“JVM Processes 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.
- **“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.
- **“Drilldowns’ Displays”** or **“Other”** displays: These displays (such as **“Alerts Table by Component”** and **“Alert Detail for Component”**) are typically accessed by drilling down from other displays (however, **“Alerts History Table”** is accessed directly from the navigation tree).

## Tomcat Displays

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”**
- **“Tomcat Servers Heatmap”**: Performance metrics for one Tomcat Server, including current and historic performance metrics.
- **“Single Tomcat Server”**: Heatmap of performance metrics for all Web modules for one Tomcat Server.
- **“All Tomcat Apps”**: Table and trend graphs of performance metrics for Web modules.
- **“Single Tomcat App”**: Table and trend graphs of performance metrics for a single Web module.

## Tomcat Overview

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:

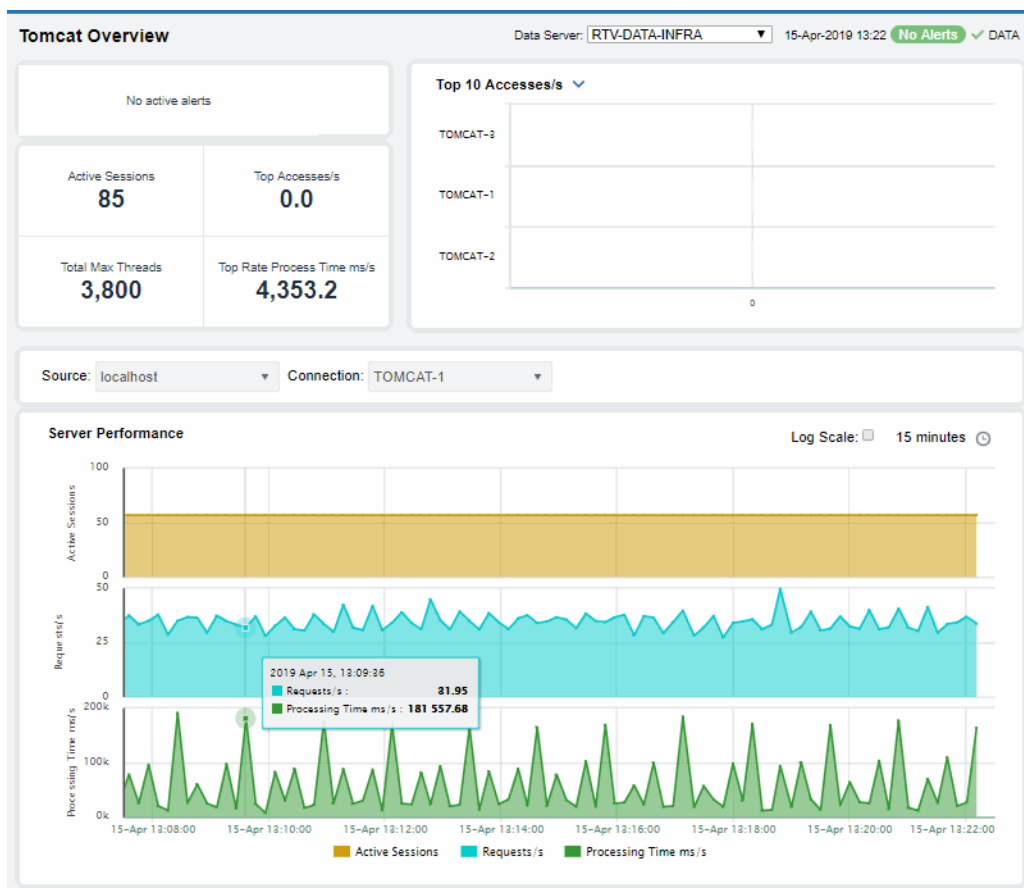
- The total number of active alerts for the selected DataServer, including the total number of critical and warning alerts.
- The greatest number of active sessions, top accesses per second, highest number of connections and the top process rate.
- A visual list of the top 10 servers with the greatest 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.

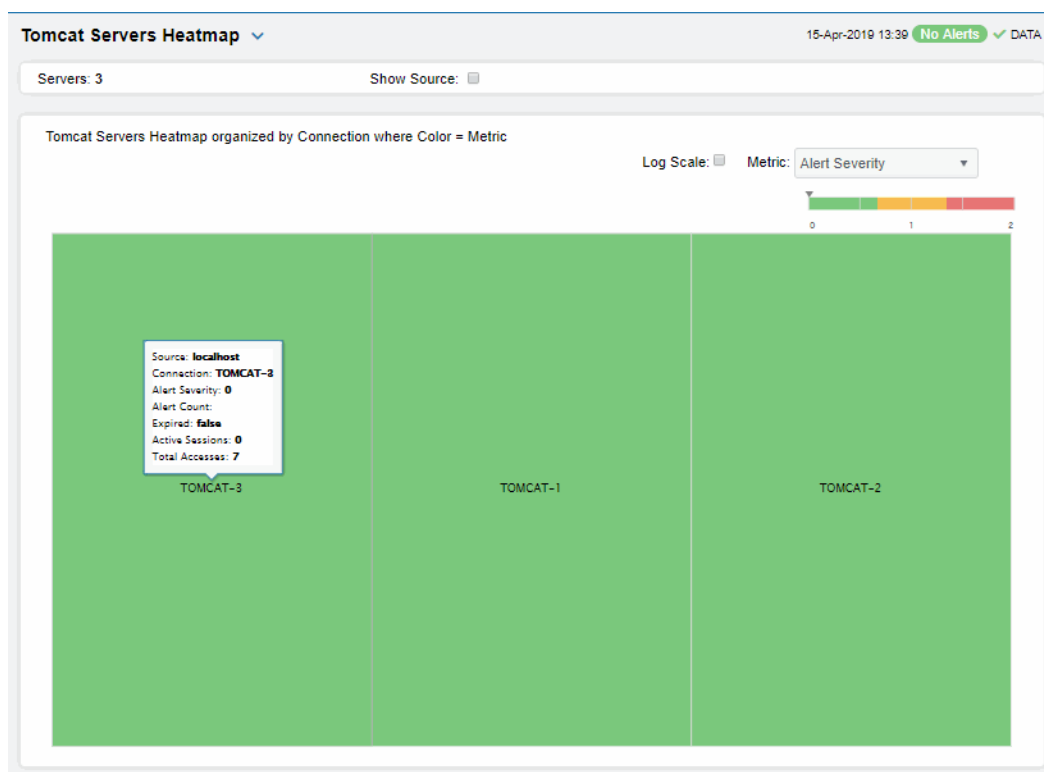


## Tomcat Servers Heatmap

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.

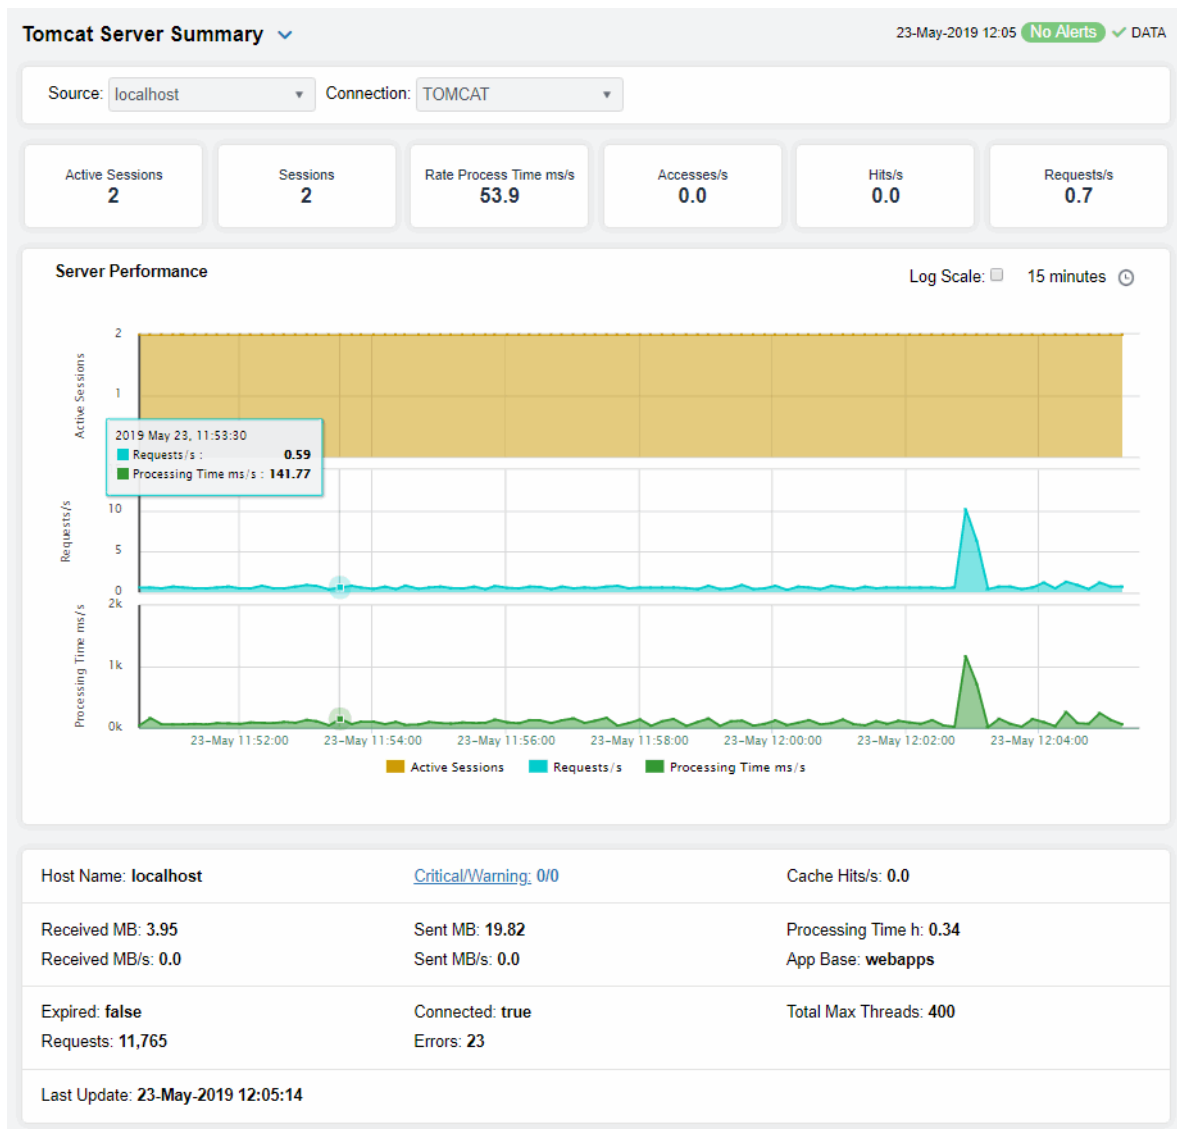


## Single Tomcat Server

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

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**.

Or just click a column header to sort.

Right-click on a table cell to **Export to Excel** or **Copy Cell Value**.

Double-click a web module to see details in the **Tomcat Application Summary display**.

| Web Module      | Expired | Alert Level | Alert Count | Total Accesses | Accesses/s | Total Sessions | Active Sessions | Expired Sessions | Cache Hits |
|-----------------|---------|-------------|-------------|----------------|------------|----------------|-----------------|------------------|------------|
| /manager        |         | ✓           |             | 7              | 0.0        | 0              | 0               | 0                | C          |
| /rtvdisplayplus |         | ✓           |             | 7              | 0.0        | 0              | 0               | 0                | C          |
| /host-manager   |         | ✓           |             | 7              | 0.0        | 0              | 0               | 0                | C          |
| /docs           |         | ✓           |             | 7              | 0.0        | 0              | 0               | 0                | C          |
| /rtvdataplus    |         | ✓           |             | 7              | 0.0        | 0              | 0               | 0                | C          |
| /rtvquery       |         | ✓           |             | 7              | 0.0        | 0              | 0               | 0                | C          |

## All Tomcat Apps

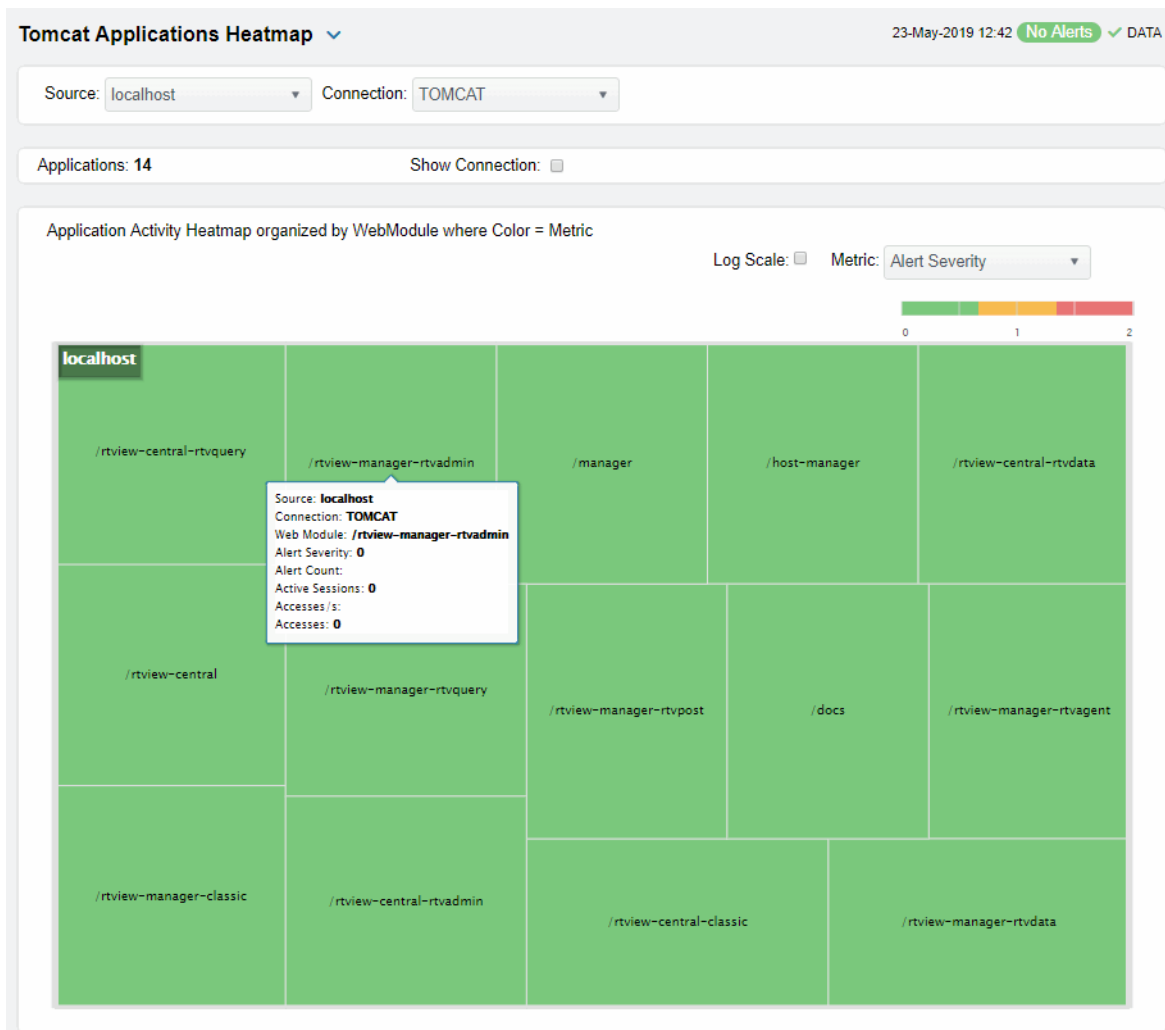
This heatmap allows you to view the status and alerts of Tomcat applications on a particular host or **All** hosts, and a particular connection or **All** connections.

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

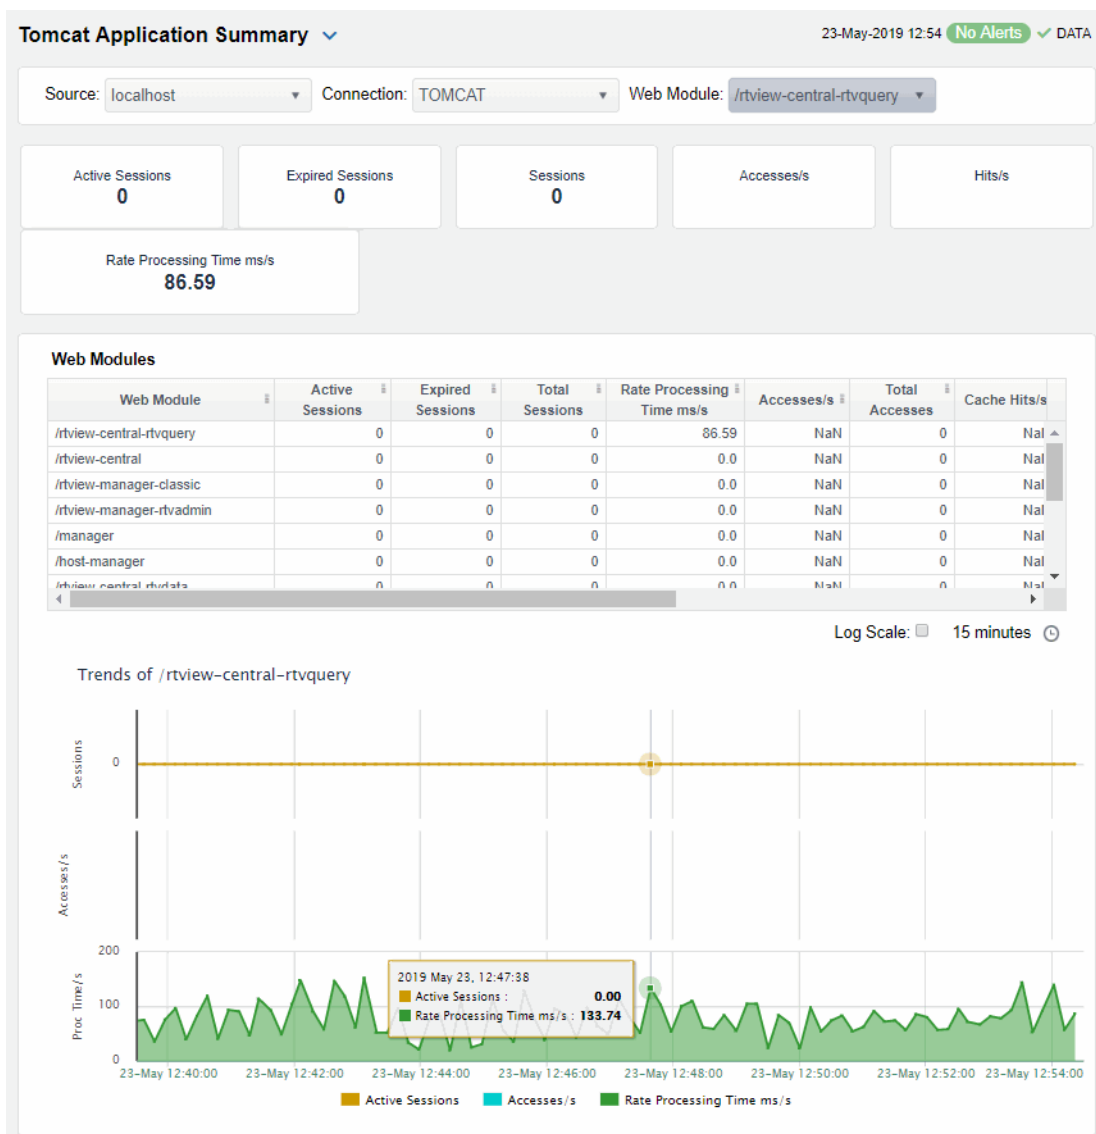
Track utilization and performance metrics for a particular Tomcat web module. Clicking on the sessions/processing rate information boxes at the top of the display takes you to the **Tomcat Servers Table** display, where you can compare and sort performance values against other Tomcat servers.

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.



## JVM Processes 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 [“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 HTML version features an overview display, [“JVM Overview”](#) (pictured below), and the following displays which can be found under **Components** tab > **Processes /JVM Processes** once RTView Manager is installed:

- [“JVMs Heatmap”](#): Heatmap of alert states for all JVM connections
- [“JVM Summary”](#): Table of connection details for all JVM connections.
- [“JVM System Properties”](#): Table of connection details for a single JVM as well as performance trend graphs.
- [“JVM GC Trends”](#): Trend graphs of garbage collection memory utilization.

### JVM Overview

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.

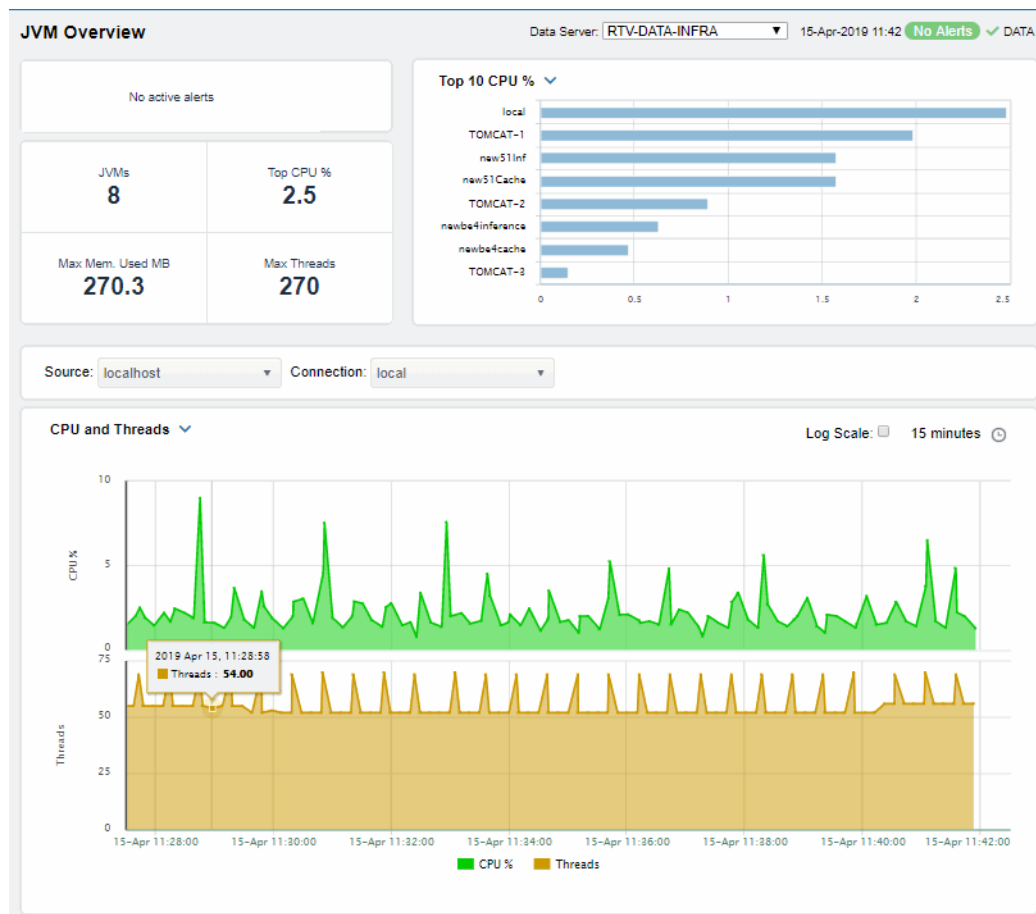
A bar graph shows the JVMs with **Top 10 CPU %** utilization. Use the drop-down menu to show JVMs with **Top 10 Used Heap** memory utilization and JVMs with **Top 10 Live Threads**.

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”](#) 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

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.

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 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.

Toggle between the commonly accessed **Table** and **Heatmap** 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**.

**JVMs Table** 15-Apr-2019 12:59 No Alerts DATA

Source: - All -

Active Only:  JVMs: 14

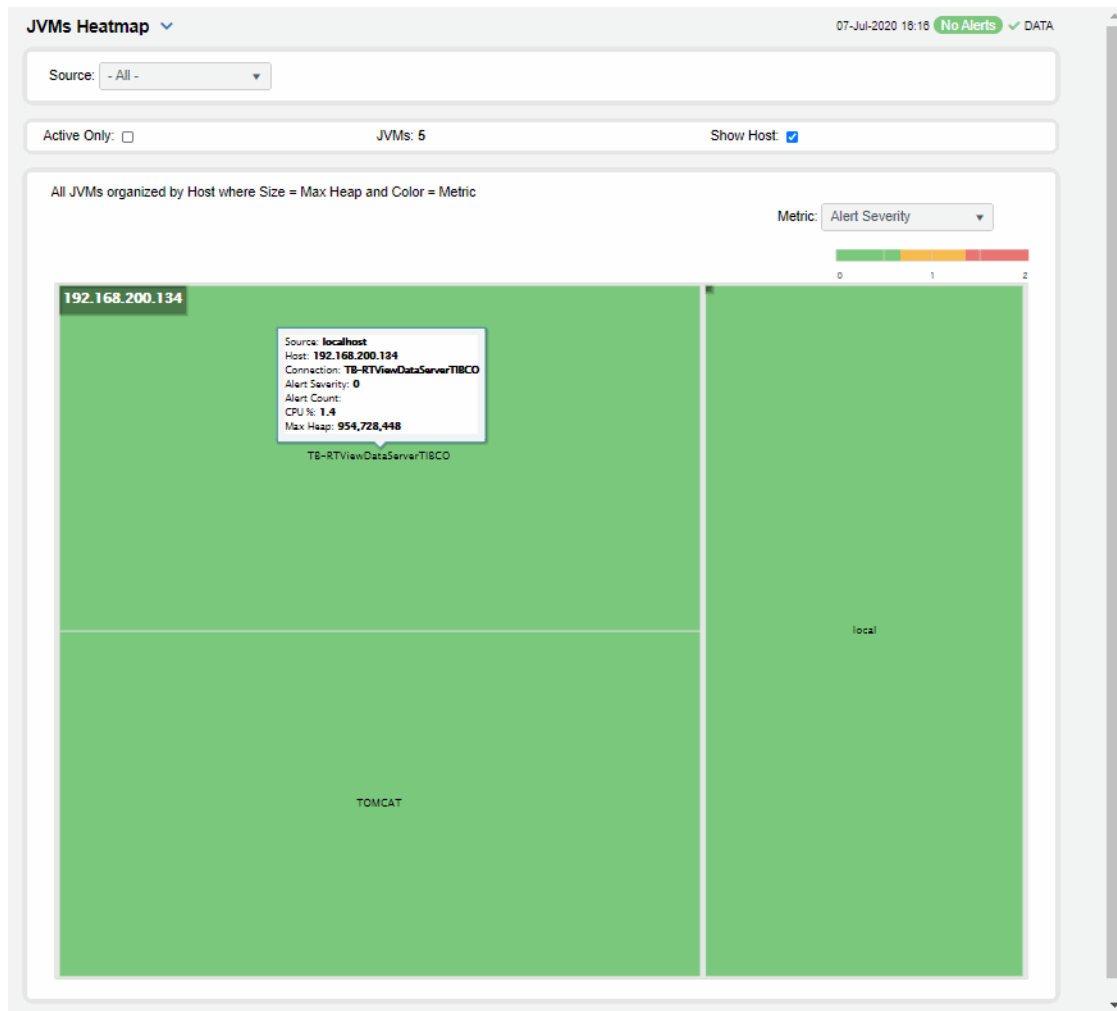
| Source    | Connection      | Expired | Connected | Alert Level | Alert Count | Host            | Port  | PID              | CPU % |
|-----------|-----------------|---------|-----------|-------------|-------------|-----------------|-------|------------------|-------|
| localhost | TOMCAT-3        |         | ✓         | ✓           |             | 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

View the 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 host, and uses color to show the selected **Metric** value for each JVM connection.

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 **Sources** 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 **PI D**). Drill-down and investigate by clicking a rectangle in the heatmap to view details for the selected connection in the **JVM Summary** display.




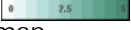


### Metric

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 0 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 0 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 0 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 0 to the maximum amount used in the heatmap.                 |

## JVM Summary

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).

Choose a **Source** and **Connection** from the drop down menus. 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**.

You can also verify whether the memory usage has reached a plateau. And if usage is approaching 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.

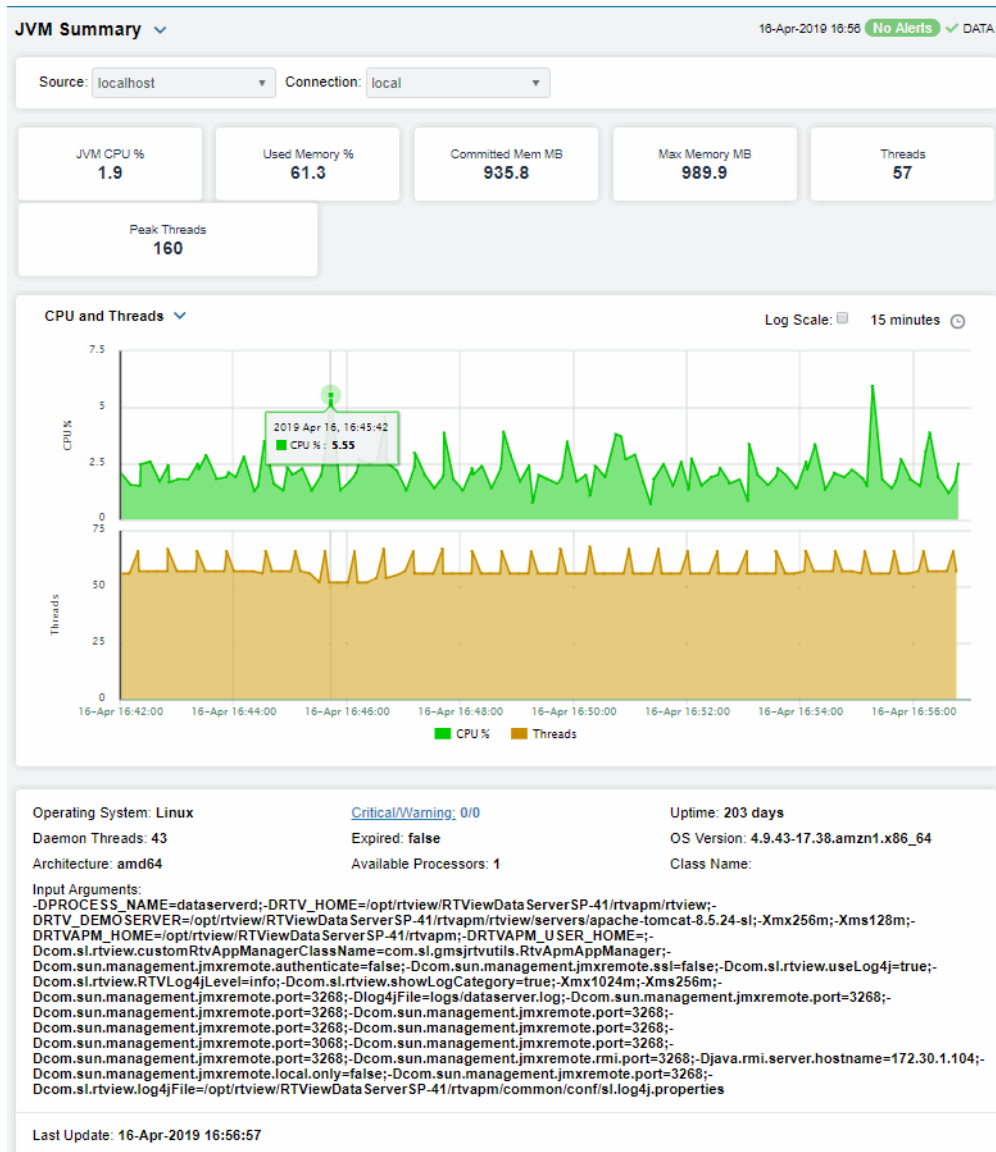
You can set the time range for the trend graph to trace. You can also choose what to trace from the drop-down menu:

- **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

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. 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**.

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.slt.rtv.customRtvAppManagerClassName=com.slt.gmsjrtvutils.RtvAppManager |
| -Dcom.sun.management.jmxremote.authenticate=false                             |
| -Dcom.sun.management.jmxremote.ssl=false                                      |
| -Dcom.slt.rtv.useLog4j=true                                                   |
| -Dcom.slt.rtv.RTVLog4jLevel=info                                              |
| -Dcom.slt.rtv.showLogCategory=true                                            |
| -Xmx:1024m                                                                    |
| -Xms256m                                                                      |

**Command-Line Arguments**

| Value |
|-------|
|-------|

**System Properties**

| Property                              | Value                    |
|---------------------------------------|--------------------------|
| com.sun.management.jmxremote.rmi.port | 3288                     |
| awt.toolkit                           | sun.awt.windows.WToolkit |
| com.slt.rtv.RTVLog4jLevel             | info                     |
| file.encoding.pkg                     | sun.io                   |
| java.specification.version            | 1.7                      |
| sun.cpu.isalist                       | amd64                    |

## JVM GC Trends

Track JVM garbage collection memory utilization trends for a single connection. Choose a **Source**, **Connection** and **Garbage Collector** 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.



## RTView Servers Displays

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"**: Shows metrics for RTView Data Servers.
- **"Data Server Summary"**: Shows details for one Data Server.
- **"Display Servers"**: Note that this display does not contain data.
- **"Display Server Summary"**: Note that this display does not contain data.
- **"Historian Servers"**: Shows metrics for RTView Historian Servers.

## Data Servers

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.

**RTView Data Servers Table** 16-Apr-2019 14:56 ✓ DATA

Source: - All -

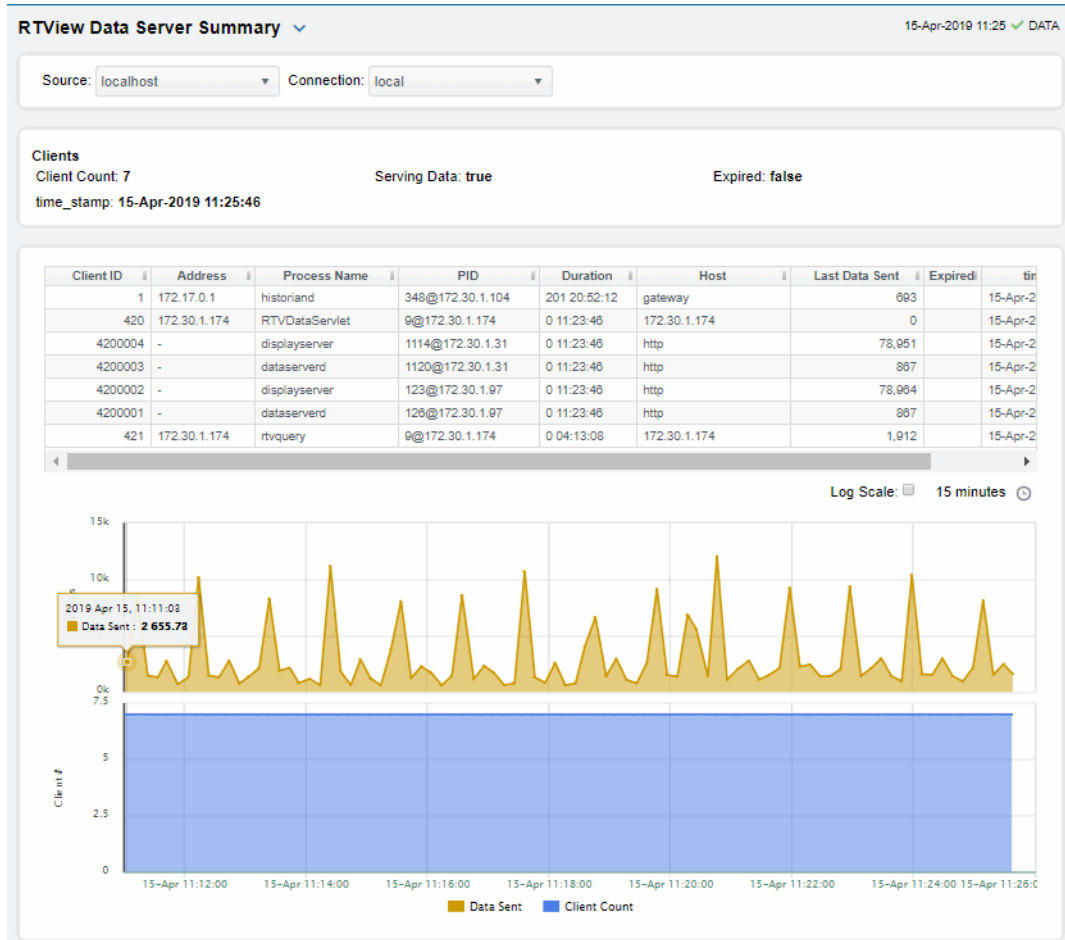
Data Servers: 1

| Source    | Connection | Client Count | Serving Data | Expired | time_stamp           |
|-----------|------------|--------------|--------------|---------|----------------------|
| localhost | local      | 7            | ✓            |         | 16-Apr-2019 14:56:16 |

## Data Server Summary

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.



## Historian Servers

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**.

| 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 |

## 'Drilldowns' Displays

This View contains the following displays:

- ["Alerts History Table"](#): Track history of any alert that has occurred in your RTView Enterprise system.
- ["Alerts Table by Component"](#): Track alerts associated with CIs shown in a display.
- ["Alert Detail for Component"](#): Investigate an alert instance and its history.

### Alerts History Table

Use this display to track the history of alerts, including cleared alerts in your monitoring system. There is one row in the table for each update to each alert.

Choose a Data Server from the drop down to filter alerts shown in the table. The **Alerts History Table** only shows alerts associated with the selected Data Server.

Select **Expand Alert Index** to separate each column in the **Alert Index** into different lines of text. When unselected, the **Alert Index** remains as a single line, with all index parts separated by semicolon (;).

Select **History Alerts** to show all historical alerts. When unselected, only current alerts are shown in the table.

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**.

Alerts History Table 07-Oct-2019 09:25 ✓ DATA

Data Server: RTV-DATA-TIBCO Expand Alert Index: ☐ History Alerts: ☑

Data Server URL: [https://rtvdemos.sl.com/tibmon\\_rtquery/](https://rtvdemos.sl.com/tibmon_rtquery/) 15 minutes

| Alert Level | Ack | Cleared | Alert Name                | Alert Index            | Alert Text                | Owner | Id     | Source | Row Update Time      |
|-------------|-----|---------|---------------------------|------------------------|---------------------------|-------|--------|--------|----------------------|
| ⚠           |     |         | JvmMemoryUsed-High        | win44-newbe4cache      | High Alert Limit exceeded |       | 144843 |        | 2019-Oct-07 09:11:41 |
| ⚠           |     | ✓       | JvmMemoryUsed-High        | win44-newbe4cache      | High Alert Limit exceeded |       | 144843 |        | 2019-Oct-07 09:11:51 |
| ⚠           |     |         | BwServerCpuUsed-High      | sl4-64(slmon)          | High Alert Limit exceeded |       | 144911 |        | 2019-Oct-07 09:18:4  |
| ⚠           |     | ✓       | BwServerCpuUsed-High      | sl4-64(slmon)          | High Alert Limit exceeded |       | 144911 |        | 2019-Oct-07 09:19:0  |
| ⚠           |     |         | BwProcessAvgElapsedTimeHk | sl4-64(slmon)-domains  | High Alert Limit exceeded |       | 144855 |        | 2019-Oct-07 09:10:5  |
| ⚠           |     |         | BwProcessAvgElapsedTimeHk | sl4-64(slmon)-domains  | High Alert Limit exceeded |       | 144854 |        | 2019-Oct-07 09:10:5  |
| ⚠           |     |         | BwProcessAvgElapsedTimeHk | sl4-64(slmon)-domains  | High Alert Limit exceeded |       | 144853 |        | 2019-Oct-07 09:10:5  |
| ⚠           |     |         | BwProcessAvgElapsedTimeHk | sl4-64(slmon)-domains  | High Alert Limit exceeded |       | 144852 |        | 2019-Oct-07 09:10:5  |
| ⚠           |     |         | BwProcessAvgElapsedTimeHk | sl4-64(slmon)-domains  | High Alert Limit exceeded |       | 144851 |        | 2019-Oct-07 09:10:5  |
| ⚠           |     |         | BwProcessAvgElapsedTimeHk | sl4-64(slmon)-domains  | High Alert Limit exceeded |       | 144856 |        | 2019-Oct-07 09:11:5  |
| ⚠           |     |         | BwProcessAvgElapsedTimeHk | slhpux11(slmon)-domain | High Alert Limit exceeded |       | 144865 |        | 2019-Oct-07 09:12:1  |
| ⚠           |     |         | BwProcessAvgElapsedTimeHk | slhpux11(slmon)-domain | High Alert Limit exceeded |       | 144864 |        | 2019-Oct-07 09:12:1  |
| ⚠           |     |         | BwProcessAvgElapsedTimeHk | slhpux11(slmon)-domain | High Alert Limit exceeded |       | 144863 |        | 2019-Oct-07 09:12:1  |
| ⚠           |     |         | BwProcessAvgElapsedTimeHk | slhpux11(slmon)-domain | High Alert Limit exceeded |       | 144862 |        | 2019-Oct-07 09:12:1  |
| ⚠           |     |         | BwProcessAvgElapsedTimeHk | slhpux11(slmon)-domain | High Alert Limit exceeded |       | 144861 |        | 2019-Oct-07 09:12:1  |
| ⚠           |     |         | BwProcessAvgElapsedTimeHk | slhpux11(slmon)-domain | High Alert Limit exceeded |       | 144860 |        | 2019-Oct-07 09:12:1  |
| ⚠           |     |         | BwProcessAvgElapsedTimeHk | slhpux11(slmon)-domain | High Alert Limit exceeded |       | 144859 |        | 2019-Oct-07 09:12:1  |
| ⚠           |     |         | BwProcessAvgElapsedTimeHk | sl4-64(slmon)-domains  | High Alert Limit exceeded |       | 144872 |        | 2019-Oct-07 09:12:1  |
| ⚠           |     |         | BwProcessAvgElapsedTimeHk | sl4-64(slmon)-domains  | High Alert Limit exceeded |       | 144871 |        | 2019-Oct-07 09:12:1  |
| ⚠           |     |         | BwProcessAvgElapsedTimeHk | sl4-64(slmon)-domains  | High Alert Limit exceeded |       | 144870 |        | 2019-Oct-07 09:12:1  |
| ⚠           |     |         | BwProcessAvgElapsedTimeHk | sl4-64(slmon)-domains  | High Alert Limit exceeded |       | 144869 |        | 2019-Oct-07 09:12:1  |
| ⚠           |     |         | BwProcessAvgElapsedTimeHk | sl4-64(slmon)-domains  | High Alert Limit exceeded |       | 144868 |        | 2019-Oct-07 09:12:1  |

### Alerts Table by Component

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 17 Alerts (the alert status icon) in the title bar of other displays. The display in which you click 17 Alerts 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:

Alert Detail 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: False      ACK: False

---

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_Lqa)   | High V |
| 2018-Oct-01 06:20:10 |             |             |         | ⚠️          | HostCpuPercentHigh     | SL-DEMO;SLHOST17(sl_Lamx)  | High A |
| 2019-May-02 03:28:5  |             |             |         | 🔔           | HostMemoryUsedHigh     | SL-DEMO-LX;192.168.200.92  | High V |
| 2018-Oct-01 06:19:38 |             |             |         | ⚠️          | HostVirtualMemoryUsedH | SL-DEMO;SLHOST17(sl_Lamx)  | High A |
| 2018-Oct-01 06:18:38 |             |             |         | 🔔           | HostMemoryUsedHigh     | SL-DEMO;SLHOST17(sl_Lamx)  | 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 10:33:5  |             |             |         | ⚠️          | HostVirtualMemoryUsedH | SL-DEMO;SLHOST16(sl_Lqa)   | 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

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  in the **Alerts Table** or  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.

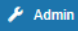


The alert reached its WARNING LEVEL threshold in the table row.


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:


 to see utilization conditions for the CI associated with the alert in a summary display.

 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.

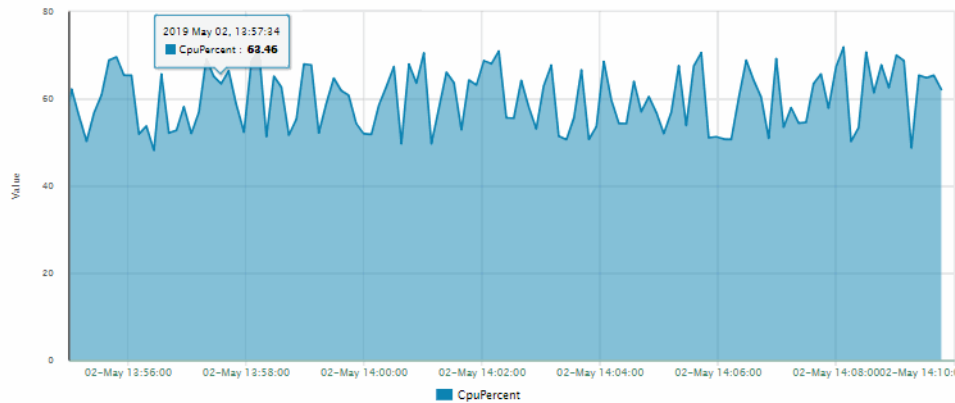
**Alert Detail For Component** ▼
02-May-2019 14:09:52 ✔ DATA OK

|                               |                                                                                             |                                                           |         |
|-------------------------------|---------------------------------------------------------------------------------------------|-----------------------------------------------------------|---------|
| Alert Name: JvmCpuPercentHigh | Severity:  | Cleared: <span style="color: green;">✔</span> 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 |  | <span style="color: green;">✔</span> | DATA UPDATE    |              |       |
| 937671 | 2019-May-02 14:03:26 |  |                                      |                |              |       |
| 937646 | 2019-May-02 13:56:27 |  | <span style="color: green;">✔</span> | DATA UPDATE    |              |       |
| 937646 | 2019-May-02 13:55:49 |  |                                      |                |              |       |
| 937635 | 2019-May-02 13:51:15 |  | <span style="color: green;">✔</span> | DATA UPDATE    |              |       |
| 937635 | 2019-May-02 13:51:14 |  |                                      |                |              |       |

## Alerts Displays

### Alerts Table

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), **Clear** to set the **Cleared** flag on the selected alert(s), **Comment** to add a comment to the alert(s) and **CI** to get details about the CI associated with the alert (these buttons are enabled when you click one or more alerts).

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 |                |  |                    |  |                 |  |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----|-----|------|--------------------|-------------------------|-------|--------|--------------|----------|-------|----------------------|------|----------------|--|--------------------|--|-----------------|--|
| <input type="button" value="Own"/> <input type="button" value="Ack"/> <input type="button" value="Unack"/> <input type="button" value="Clear"/> <input type="button" value="Comment"/> <input type="button" value="Details"/> <input type="button" value="CI"/> |     |     |      |                    |                         |       |        |              |          |       | 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

---

## Admin Displays

These displays enable you to set alert thresholds, observe how alerts are managed, and view internal data gathered and stored by RTView (used for troubleshooting with SL Technical Support). Displays in this View are:

- **“Alert Administration”**: Displays active alerts and provides interface to modify, enable and manage alerts.
- **“Alert Overrides Admin”**: Set and modify alert overrides. Access this display from the **Alert Administration** display.
- **“Cache Table”**: View cached data that RTView is capturing and maintaining, and use this data use this for debugging with SL Technical Support.

### Alert Administration

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.

You can also override the alert duration time per alert index instead of to all indexes. To override the duration for an alert index, select the alert in the **Alert Administration** display, click **Override** and edit the **Alert Delay**. For alert indexes that were overridden in a previous release (before duration override was supported) the override duration is set to **-1**, indicating that this is set to use the top level alert duration.

### 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.

The screenshot shows the 'Alerts Administration' interface. At the top, it displays the date '30-Apr-2019 10:34:01' and a status 'DATA OK'. Below this is a 'Package' dropdown set to 'All' and a URL 'http://rtvdemos.sl.com/emdemo\_central\_rtquery'. The main part of the interface is a table with the following columns: Alert Name, Alert Enabled, Alert Delay, Warning Level, Alert Level, and Override Count. The table lists 14 alerts, with 'JvmThreadCountHigh' selected. Below the table is a 'Settings for alert' section for the selected alert, 'HostSwapUsedHigh'. This section includes fields for 'Alert Enabled' (unchecked), 'Delay' (30), 'Warning Level' (75), and 'Alert Level' (90). There are three buttons: 'Save Settings', 'Original Defaults', and 'Override Settings'. At the bottom, it shows 'Alert Selected: HostSwapUsedHigh' and its description: 'The percentage of swap space used is above the limits defined for that Host'.

| 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: HostSwapUsedHigh  
 Alert Enabled:  Delay: 30 Warning Level: 75 Alert Level: 90  
 Save Settings Original Defaults Override Settings  
 Alert Selected: HostSwapUsedHigh Description: The percentage of swap space used is above the limits defined for that Host

For additional details, see ["Alert Overrides Admin"](#).

|                      |                                                                                                                                                                                     |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Alert Name</b>    | The name of the alert.                                                                                                                                                              |
| <b>Alert Enabled</b> | When checked, the alert is enabled globally.                                                                                                                                        |
| <b>Alert Delay</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>Warning Level</b> | The global warning threshold for the selected alert. When the specified value is exceeded a warning is executed.                                                                    |

|                       |                                                                                                                                                                                                                                                                     |
|-----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Alert Level</b>    | The global alarm threshold for the selected alert. When the specified value is exceeded an alarm is executed.                                                                                                                                                       |
| <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:<br>-0 indicates that no overrides are applied to the alert.<br>-1 indicates that the alert does not support overrides. |

#### Settings for alert

Select an alert in the table to use the following options:

|                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Alert Enabled</b>     | Check / uncheck this box to enable or disable the selected alert globally.                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Delay</b>             | Enter the amount of time (in seconds) that the value must be above the specified Warning Level or Alarm Level threshold before the selected alert is executed. 0 is for immediate execution.                                                                                                                                                                                                                                                                             |
| <b>Warning Level</b>     | Enter 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.                                                                                                                                                                                                          |
| <b>Alert Level</b>       | Enter 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.<br>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>Save Settings</b>     | Click to apply alert settings for the selected alert.                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Original Defaults</b> | Click to revert to original alert settings for the selected alert.                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Override Settings</b> | Click to set an alert override in the <b>Alert Overrides Admin</b> display on the selected alert.                                                                                                                                                                                                                                                                                                                                                                        |

## Alert Overrides Admin

Administrators use this display to override the alert settings defined in the **Alert Administration** display. To access this display, select an alert in the **Alert Administration** display and choose **Override Settings**.

The table lists all the possible overrides that can be defined for the alert you selected from the **Alert Administration** display. Each row in the table represents a different resource or group of resources that can be overridden. When the four last columns are blank, that means the resource has not been overridden, and the default settings for the alert apply. Otherwise, columns describe whether the alert is enabled, if the override itself is enabled, the overridden alert thresholds and the overridden duration for each row.

Use the **Override Type** drop-down menu to switch the list to a specific type of override (the options for this menu vary according to the alert type), and use the **Display** drop-down menu to list **All** resources, **Overridden** resources or **Free** resources.

You can also enter a pattern or regular expression in the **Search** string to limit the list.

The **RegEx** checkbox indicates whether the text you entered is treated as a search pattern or as a regular expression. Multiple rows can be selected to create/edit/remove many overrides simultaneously.

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 overrides:

Select an **Override Type** from the drop-down menu (depending on the alert, there might be only one type) and then select one or more rows from the table. Under **Settings for selected index** (in the lower portion of the screen), modify settings for the **Override Enabled**, **Alert Enabled**, **Alert Delay**, **Warning Level** and/or **Alarm Level**, then click **Add Override**. The table updates with your new settings.

#### To alter overrides:

To alter existing overrides with new settings, select them from the table, set all properties under **Settings for selected index** as desired, then click **Save Settings**. To clear existing overrides, select one or more rows, then click **Remove Override**.

**Note:** You can override alert and warning levels without overriding duration by setting it to **-1**.

For alert indexes that were overridden in a previous release (before duration override was supported) the override duration is set to **-1**, indicating that this is set to use the top level alert duration.

**Alert Overrides Administration** Data Server: TB-DataServerInfra 10-Jun-2020 16:01:28 ✓DATA OK

Alert: TomcatAppAccessRateHigh Override Type: PerApplication Display: All

Search: \* RegEx:

| Source    | Connection | path              | Override Enab | Alert Enabled | Alert Delay | Warning Level | Alert Level |
|-----------|------------|-------------------|---------------|---------------|-------------|---------------|-------------|
| localhost | TOMCAT     | /rtview-central-r |               |               |             |               |             |
| localhost | TOMCAT     | /rtview-manage    |               |               |             |               |             |
| localhost | TOMCAT     | /rtview-central   |               |               |             |               |             |
| localhost | TOMCAT     | /rtview-manage    |               |               |             |               |             |
| localhost | TOMCAT     | /manager          |               |               |             |               |             |
| localhost | TOMCAT     | /host-manager     |               |               |             |               |             |
| localhost | TOMCAT     | /rtview-central-r |               |               |             |               |             |

**Settings for selected index**

Override Enabled:  Alert Enabled:  Alert Delay: 30 Warning Level: 50

Alert Level: 100

**Add Override** **Save Settings** **Remove Override**

## Cache Table

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:  History Tables:

Data Server URL: [https://rtvdemos.sl.com/emdemo\\_central\\_rtquery](https://rtvdemos.sl.com/emdemo_central_rtquery)

| 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   | 8       | 12506  |

Cache: RtvAlertStatsByCIAndAlertGroup Table: current

| time_stamp           | CITYPE | CINAME                  | ALERTGROUP | MaxSeverity | AlertCount |
|----------------------|--------|-------------------------|------------|-------------|------------|
| 2019-May-07 14:11:33 | JVM    | localhost:SOLMON_TOM    | None       | 2           | 1          |
| 2019-May-07 14:11:33 | JVM    | localhost:EMSMON_TOI    | 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_DATI   | 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;CRMZooko  | 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

## Alerts for RTView Manager

RTView Manager comes with the following alert types for RTView Servers (Data Servers, Display Servers and Historian Servers):

### JvmCpuPercentHigh

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

### JvmGcDutyCycleHigh

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.<br>Index Type: Per GC Source<br>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.<br>Index Type(s): Per JVM<br>Metric: MemoryUsedPercent                                |
| <b>JvmNotConnected</b>             | Executes a single alert if the JVM is disconnected, indicating that it might have crashed.<br>Index Type(s): Per JVM<br>Metric: Connected                                                                     |
| <b>JvmStaleData</b>                | Executes a single alert if the data update wait time exceeds the specified duration threshold.<br>Index Type(s): Per JVM<br>Metric: Expired                                                                   |
| <b>JvmThreadCountHigh</b>          | Executes a single warning alert and a single alarm alert if the number of threads exceeds the specified threshold.<br>Index Type(s): Per JVM<br>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.<br>Index Type(s): Per Server<br>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.<br>Index Type(s): Per Server<br>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.<br>Index Type(s): Per Application<br>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.<br>Index Type(s): Per Application<br>Metric: activeSessions                        |

---

## Configure High Availability

To configure HA for RTView Manager, refer to instructions in the *RTView Enterprise User's Guide*.



## 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"
- "IBM MQ"
- "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>AcwI nstanceCpuHigh</b><br>Executes a single warning and single alarm if the CPU used exceeds the specified threshold.<br>Index Type: PerInstance<br>Metric: CPUUtilization                                             | 70            | 80          | 30       | FALSE   |
| <b>AcwI nstanceDiskReadBytesHigh</b><br>Executes a single warning and single alarm if the number of bytes read from the disk exceeds the specified threshold.<br>Index Type: PerInstance<br>Metric: DiskReadBytes          | 10000         | 20000       | 30       | FALSE   |
| <b>AcwI nstanceDiskReadOpsHigh</b><br>Executes a single warning and single alarm if the number of disk reads exceeds the specified threshold.<br>Index Type: PerInstance<br>Metric: DiskReadOps                            | 100           | 200         | 30       | FALSE   |
| <b>AcwI nstanceDiskWriteBytesHigh</b><br>Executes a single warning and single alarm if the number of bytes written to the disk exceeds the specified threshold.<br>Index Type: PerInstance<br>Metric: DiskWriteBytes       | 1000000       | 2000000     | 30       | FALSE   |
| <b>AcwI nstanceDiskWriteOpsHigh</b><br>Executes a single warning and single alarm if the number of disk writes exceeds the specified threshold.<br>Index Type: PerInstance<br>Metric: DiskWriteOps                         | 100           | 200         | 30       | FALSE   |
| <b>AcwI nstanceNetworkReadBytesHigh</b><br>Executes a single warning and single alarm if the number of bytes read from the network exceeds the specified threshold.<br>Index Type: PerInstance<br>Metric: NetworkIn        | 1000000       | 20000       | 30       | FALSE   |
| <b>AcwI nstanceNetworkWriteBytesHigh</b><br>Executes a single warning and single alarm if the number of bytes written across the network exceeds the specified threshold.<br>Index Type: PerInstance<br>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.<br>LEVEL | ALARM<br>LEVEL | DURATION | ENABLED |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------|----------------|----------|---------|
| <b>KafkaBrokerBytesInPerSecHigh</b><br>The number of incoming bytes per second exceeds the defined threshold for the broker.<br><b>Index Type(s):</b> PerKafkaServer                                      | 1600           | 2000           | 30       | FALSE   |
| <b>KafkaBrokerBytesOutPerSecHigh</b><br>The number of outgoing bytes per second exceeds the defined threshold for the broker.<br><b>Index Type(s):</b> PerKafkaServer                                     | 1600           | 2000           | 30       | FALSE   |
| <b>KafkaBrokerCpuPercentHigh</b><br>The CPU percentage reported by the JVM is above the limits defined for that broker.<br><b>Index Type(s):</b> PerKafkaServer                                           | 50             | 75             | 30       | FALSE   |
| <b>KafkaBrokerExpired</b><br>The Kafka Broker is not responding.<br><b>Index Type(s):</b> PerKafkaServer                                                                                                  | NaN            | NaN            | 30       | FALSE   |
| <b>KafkaBrokerFetchRequestsPerSecHigh</b><br>Fetch requests per second exceeds threshold for the broker.<br><b>Index Type(s):</b> PerKafkaServer                                                          | 1600           | 2000           | 30       | FALSE   |
| <b>KafkaBrokerLogFlushLatency95PHigh</b><br>The current log flush latency exceeds the 95th percentile.<br><b>Index Type(s):</b> PerKafkaServer                                                            | 1600           | 2000           | 30       | FALSE   |
| <b>KafkaBrokerMemoryUsedPercentHigh</b><br>The percentage of heap memory used relative to the maximum heap available is above the limits defined for that broker.<br><b>Index Type(s):</b> PerKafkaServer | 50             | 75             | 30       | FALSE   |
| <b>KafkaBrokerMsgsInPerSecHigh</b><br>The number of incoming messages per second exceeds the defined threshold for the broker.<br><b>Index Type(s):</b> PerKafkaServer                                    | 1600           | 2000           | 30       | TRUE    |
| <b>KafkaBrokerNetProcAvgIdlePctHigh</b><br>The average percent idle for the network processor exceeds the threshold.<br><b>Index Type(s):</b> PerKafkaServer                                              | 1              | .3             | 30       | FALSE   |

|                                                                                                                                          |      |      |    |       |
|------------------------------------------------------------------------------------------------------------------------------------------|------|------|----|-------|
| <b>KafkaBrokerNetProcAvgIdlePctLow</b>                                                                                                   | .05  | .3   | 30 | FALSE |
| The average percent idle for the network processor is below the threshold.<br><b>Index Type(s):</b> PerKafkaServer                       |      |      |    |       |
| <b>KafkaBrokerOfflinePartitionCountHigh</b>                                                                                              | NaN  | 1    | 30 | TRUE  |
| The number of partitions without an active leader is not zero.<br><b>Index Type(s):</b> PerKafkaServer                                   |      |      |    |       |
| <b>KafkaBrokerProduceRequestsPerSecHigh</b>                                                                                              | 1600 | 2000 | 30 | FALSE |
| Produce requests per second exceeds threshold for the broker.<br><b>Index Type(s):</b> PerKafkaServer                                    |      |      |    |       |
| <b>KafkaBrokerUnclnLderElecsPerSecHigh</b>                                                                                               | 1600 | 2000 | 30 | FALSE |
| The available replicas were not in sync during leader election. Data loss has probably occurred.<br><b>Index Type(s):</b> PerKafkaServer |      |      |    |       |
| <b>KafkaBrokerUnderReplicatedPartnsHigh</b>                                                                                              | NaN  | 1    | 30 | FALSE |
| The number of under-replicated partitions is not zero.<br><b>Index Type(s):</b> PerKafkaServer                                           |      |      |    |       |
| <b>KafkaClusterLeadersUnbalancedHigh</b>                                                                                                 | 10   | 10   | 30 | FALSE |
| The partition leaders for the cluster are not evenly distributed across the available brokers.<br><b>Index Type(s):</b> PerKafkaCluster  |      |      |    |       |
| <b>KafkaClusterNoActiveController</b>                                                                                                    | NaN  | NaN  | 30 | FALSE |
| There is more than one active controller per cluster, which could indicate a split-brain error.<br><b>Index Type(s):</b> PerKafkaCluster |      |      |    |       |
| <b>KafkaClusterPartitionsUnbalancedHigh</b>                                                                                              | 10   | 1    | 30 | FALSE |
| Partitions supported by the cluster are not evenly distributed across the available brokers.<br><b>Index Type(s):</b> PerKafkaCluster    |      |      |    |       |
| <b>KafkaClusterSplitBrain</b>                                                                                                            | NaN  | NaN  | 30 | FALSE |
| One (or more) zookeeper/broker is not acting as part of the main cluster.<br><b>Index Type(s):</b> PerKafkaCluster                       |      |      |    |       |
| <b>KafkaClusterTopicReplicasOutOfSync</b>                                                                                                | 5    | 10   | 30 | FALSE |
| Topic partition replicas are out of sync.<br><b>Index Type(s):</b> PerKafkaCluster                                                       |      |      |    |       |
| <b>KafkaoConsumerBytesPerSecHigh</b>                                                                                                     | 1600 | 2000 | 30 | FALSE |
| The consumer message load (bytes per second) exceeds the threshold.<br><b>Index Type(s):</b> PerKafkaConsumer                            |      |      |    |       |
| <b>KafkaoConsumerCpuPercentHigh</b>                                                                                                      | 50   | 75   | 30 | FALSE |
| The CPU percentage reported by the JVM is above the limits defined for that consumer.<br><b>Index Type(s):</b> PerKafkaConsumer          |      |      |    |       |



|                                                                                                                                                                                                                                                                                                                                  |      |      |     |       |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------|------|-----|-------|
| <b>KafkaConsumerExpired</b><br>The consumer is not responding.<br><b>Index Type(s):</b> PerKafkaConsumer                                                                                                                                                                                                                         | NaN  | NaN  | 30  | FALSE |
| <b>KafkaConsumerFetchLatencyHigh</b><br>The consumer fetch latency exceeds the threshold.<br><b>Index Type(s):</b> PerKafkaConsumer                                                                                                                                                                                              | 1600 | 2000 | 30  | TRUE  |
| <b>KafkaConsumerFetchRateHigh</b><br>The consumer is pulling records from Kafka at a slower than expected rate.<br><b>Index Type(s):</b> PerKafkaConsumer                                                                                                                                                                        | 1600 | 2000 | 30  | FALSE |
| <b>KafkaConsumerLagHigh</b><br>The consumer is falling too far behind the producer.<br><b>Index Type(s):</b> PerKafkaConsumer                                                                                                                                                                                                    | 1600 | 2000 | 30  | TRUE  |
| <b>KafkaConsumerLagIncreasing</b><br>The consumer lag rate of change is greater than zero for the specified duration, which could mean that lag is steadily increasing.<br><b>Index Type(s):</b> PerKafkaConsumer                                                                                                                | NaN  | NaN  | 300 | FALSE |
| <b>KafkaConsumerMemoryUsedPercentHigh</b><br>The percentage of heap memory used relative to the maximum heap available is above the limits defined for that consumer.<br><b>Index Type(s):</b> PerKafkaConsumer                                                                                                                  | 50   | 75   | 300 | FALSE |
| <b>KafkaConsumerPartitionStalled</b><br>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.<br><b>Index Type(s):</b> PerKafkaConsumer | NaN  | NaN  | 300 | FALSE |
| <b>KafkaConsumerRecordsConsumedRateHigh</b><br>The consumer message load (messages per second) exceeds the threshold.<br><b>Index Type(s):</b> PerKafkaConsumer                                                                                                                                                                  | 1600 | 2000 | 30  | TRUE  |
| <b>KafkaConsumerSlow</b><br>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.<br><b>Index Type(s):</b> PerKafkaConsumer                                           | NaN  | NaN  | 300 | FALSE |
| <b>KafkaProducerCpuPercentHigh</b><br>The CPU percentage reported by the JVM is above the limits defined for that producer.<br><b>Index Type(s):</b> PerKafkaProducer                                                                                                                                                            | 50   | 75   | 30  | FALSE |
| <b>KafkaProducerExpired</b><br>The producer is not responding.<br><b>Index Type(s):</b> PerKafkaProducer                                                                                                                                                                                                                         | NaN  | NaN  | 30  | FALSE |

|                                                                                                                                                                                                                    |      |      |    |       |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------|------|----|-------|
| <b>KafkaProducerIncomingByteRateHigh</b><br>The producer's incoming byte rate exceeds the threshold.<br><b>Index Type(s):</b> PerKafkaProducer                                                                     | 1600 | 2000 | 30 | TRUE  |
| <b>KafkaProducerIoWaitTimeMSHigh</b><br>The producer is waiting for IO longer than expected (on average).<br><b>Index Type(s):</b> PerKafkaProducer                                                                | 1600 | 2000 | 30 | FALSE |
| <b>KafkaProducerMemoryUsedPercentHigh</b><br>The percentage of heap memory used relative to the maximum heap available is above the limits defined for that producer.<br><b>Index Type(s):</b> PerKafkaProducer    | 50   | 75   | 30 | FALSE |
| <b>KafkaProducerOutgoingByteRateHigh</b><br>The producer output byte rate exceeds the threshold.<br><b>Index Type(s):</b> PerKafkaProducer                                                                         | 1600 | 2000 | 30 | TRUE  |
| <b>KafkaProducerRecordSendRateHigh</b><br>The producer record send rate exceeds the threshold.<br><b>Index Type(s):</b> PerKafkaProducer                                                                           | 1600 | 2000 | 30 | TRUE  |
| <b>KafkaProducerRequestLatencyHigh</b><br>The producer request latency exceeds the threshold.<br><b>Index Type(s):</b> PerKafkaProducer                                                                            | 1600 | 2000 | 30 | TRUE  |
| <b>KafkaProducerRequestRateHigh</b><br>The producers request rate exceeds the threshold.<br><b>Index Type(s):</b> PerKafkaProducer                                                                                 | 1600 | 2000 | 30 | TRUE  |
| <b>KafkaProducerResponseRateHigh</b><br>The producer response rate exceeds the threshold.<br><b>Index Type(s):</b> PerKafkaProducer                                                                                | 1600 | 2000 | 30 | TRUE  |
| <b>KafkaZookeeperAvgLatencyHigh</b><br>The average time for the zookeeper to respond to a request exceeds the threshold.<br><b>Index Type(s):</b> PerKafkaZookeeper                                                | 1600 | 2000 | 30 | TRUE  |
| <b>KafkaZookeeperCpuPercentHigh</b><br>The CPU percentage reported by the JVM is above the limits defined for that zookeeper.<br><b>Index Type(s):</b> PerKafkaZookeeper                                           | 50   | 75   | 30 | TRUE  |
| <b>KafkaZookeeperExpired</b><br>The zookeeper is not responding.<br><b>Index Type(s):</b> PerKafkaZookeeper                                                                                                        | NaN  | NaN  | 30 | FALSE |
| <b>KafkaZookeeperMemoryUsedPercentHigh</b><br>The percentage of heap memory used relative to the maximum heap available is above the limits defined for that Zookeeper.<br><b>Index Type(s):</b> PerKafkaZookeeper | 50   | 75   | 30 | FALSE |

|                                                                                                                                                                  |      |      |    |       |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------|------|------|----|-------|
| <b>KafkaZookeeperNumAliveConnsHigh</b><br>The total number of connections to a given zookeeper exceeds the threshold.<br><b>Index Type(s):</b> PerKafkaZookeeper | 1600 | 2000 | 30 | TRUE  |
| <b>KafkaZookeeperOutstandingReqsHigh</b><br>Clients are making requests faster than the zookeeper can process them.<br><b>Index Type(s):</b> PerKafkaZookeeper   | 1600 | 2000 | 30 | FALSE |
| <b>KafkaZookeeperRatePktsRcvdHigh</b><br>The rate that the zookeeper is receiving packets exceeds the threshold.<br><b>Index Type(s):</b> PerKafkaZookeeper      | 1600 | 2000 | 30 | TRUE  |
| <b>KafkaZookeeperRatePktsSentHigh</b><br>The rate that the zookeeper is sending packets exceeds the threshold.<br><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><br>A Docker Container's CPU usage is above the defined threshold.<br><b>Index Type(s):</b> PerContainer<br><b>Metric:</b> cpu.usage                            | 24          | 50          | 30       | FALSE   |
| <b>DocContainerExpired</b><br>A Docker Container has expired.<br><b>Index Type(s):</b> PerContainer<br><b>Metric:</b> Expired                                                                  | NaN         | NaN         | 30       | FALSE   |
| <b>DocContainerNetBytesInHigh</b><br>A Docker Container's incoming network data rate is above the defined thresholds.<br><b>Index Type(s):</b> PerContainer<br><b>Metric:</b> net.rxbytes.avg  | 750000      | 1000000     | 30       | FALSE   |
| <b>DocContainerNetBytesOutHigh</b><br>A Docker Container's outgoing network data rate is above the defined thresholds.<br><b>Index Type(s):</b> PerContainer<br><b>Metric:</b> net.txbytes.avg | 750000      | 1000000     | 30       | FALSE   |

|                                                                                                                                                                                       |        |         |    |       |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------|---------|----|-------|
| <b>DocEngineCpuUsageHigh</b><br>A Docker Engine's CPU usage is above the defined thresholds.<br><b>Index Type(s):</b> PerEngine<br><b>Metric:</b> cpu.usage                           | 50     | 75      | 30 | TRUE  |
| <b>DocEngineExpired</b><br>A Docker Engine has expired.<br><b>Index Type(s):</b> PerEngine<br><b>Metric:</b> Expired                                                                  | NaN    | NaN     | 30 | FALSE |
| <b>DocEngineNetBytesInHigh</b><br>A Docker Engine's incoming network data rate is above the defined thresholds.<br><b>Index Type(s):</b> PerEngine<br><b>Metric:</b> net.rxbytes.avg  | 750000 | 1000000 | 30 | TRUE  |
| <b>DocEngineNetBytesOutHigh</b><br>A Docker Engine's outgoing network data rate is above the defined thresholds.<br><b>Index Type(s):</b> PerEngine<br><b>Metric:</b> net.txbytes.avg | 750000 | 1000000 | 30 | TRUE  |

## IBM MQ

The following alerts are available for IBM MQ. Default settings for warning and alarm thresholds, duration and whether the alert is enabled (true/false) are shown.

| Alert                                                                                                                                                                               | Warning Level | Alarm Level | Delay | Enabled |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|-------------|-------|---------|
| <b>MqBrokerQueueDepthHigh</b><br>Executes a single warning and single alarm if the broker queue depth exceeds the specified threshold.<br>Index Type: All                           | 3000          | 4000        | 30    | FALSE   |
| <b>MqChannelBytesInHigh</b><br>Executes a single warning and single alarm if the number of bytes received by the channel exceeds the specified threshold.<br>Index Type: PerChannel | 6000          | 8000        | 30    | FALSE   |
| <b>MqChannelBytesOutHigh</b><br>Executes a single warning and single alarm if the number of bytes sent by the channel exceeds the specified threshold.<br>Index Type: PerChannel    | 6000          | 8000        | 30    | FALSE   |

|                                                                                                                                                                                               |      |      |    |       |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------|------|----|-------|
| <b>MqChannelInactive</b><br>This alert will be executed when non-system channels have the Status metric set to a value different than RUNNING.<br>Index Type: PerChannel                      | NaN  | NaN  | 30 | FALSE |
| <b>MqManagedQueueNoActiveReaders</b><br>The managed queue does not have any readers.<br>Index Type: PerQueue                                                                                  | NaN  | NaN  | 30 | FALSE |
| <b>MqManagedQueueTooManyReaders</b><br>The number of active readers for the managed queue equals or exceeds the specified threshold<br>Index Type: PerQueue                                   | NaN  | 2    | 30 | FALSE |
| <b>MqQueueDepthHigh</b><br>Executes a single warning and single alarm if the queue depth exceeds the specified threshold.<br>Index Type: PerQueue                                             | 3000 | 4000 | 30 | FALSE |
| <b>MqQueueFull</b><br>The queue has reached maximum capacity.<br>Index Type: PerQueue                                                                                                         | NaN  | NaN  | 30 | FALSE |
| <b>MqQueueMgrConnectionCountHigh</b><br>Executes a single warning and single alarm if the number of connections in the queue manager exceeds the specified threshold.<br>Index Type: PerQueue | 6    | 8    | 30 | FALSE |
| <b>MqQueueMgrDisconnected</b><br>Executes a single alert when the data server stops getting information from the queue manager.<br>Index Type: PerQueue                                       | NaN  | NaN  | 30 | FALSE |
| <b>MqQueueMgrNotRunning</b><br>Executes a single alert when the queue manager reports a status other than "RUNNING".<br>Index Type: PerQueue                                                  | NaN  | NaN  | 30 | FALSE |
| <b>MqQueueOpenInputCountHigh</b><br>Executes a single warning and single alarm if the open input count of the queue exceeds the specified threshold.<br>Index Type: PerQueue                  | 20   | 40   | 30 | FALSE |
| <b>MqQueueOpenInOutputCountHigh</b><br>Executes a single warning and single alarm if the open output count of the queue exceeds the specified threshold.<br>Index Type: PerQueue              | 20   | 40   | 30 | FALSE |
| <b>MqQueuePercentHigh</b><br>Executes a single warning and single alarm if the percent of the queue used exceeds the specified threshold.<br>Index Type: PerQueue                             | 60   | 80   | 30 | FALSE |

|                                                                                                                                                                 |      |      |    |       |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|------|------|----|-------|
| <b>MqQueueStatusNotRunning</b><br>The status of the queue indicates that it is not running.<br>Index Type: PerQueue                                             | NaN  | NaN  | 30 | FALSE |
| <b>MqSubscriptionExpired</b><br>The subscription has expired.<br>Index Type: PerMqSubscription                                                                  | NaN  | NaN  | 30 | FALSE |
| <b>MqSubscriptionMessageRateHigh</b><br>The message rate exceeds the threshold for the subscription.<br>Index Type: PerMqSubscription                           | 1600 | 2000 | 30 | FALSE |
| <b>MqSubscriptionMessageRateLow</b><br>The message rate is below the threshold for the subscription.<br>Index Type: PerMqSubscription                           | 15   | 5    | 30 | FALSE |
| <b>MqTopicExpired</b><br>The topic has expired.<br>Index Type: PerMqTopic                                                                                       | NaN  | NaN  | 30 | FALSE |
| <b>MqTopicPublishCountHigh</b><br>The publish count exceeds the threshold for the topic.<br>Index Type: PerMqTopic                                              | 1600 | 2000 | 30 | FALSE |
| <b>MqTopicPublishCountLow</b><br>The publish count is below the threshold for the topic.<br>Index Type: PerMqTopic                                              | 15   | 5    | 30 | FALSE |
| <b>MqTopicPublisherExpired</b><br>The topic publisher has expired.<br>Index Type: PerMqTopicPublisher                                                           | NaN  | NaN  | 30 | FALSE |
| <b>MqTopicPublisherPublishRateHigh</b><br>The publish rate exceeds the threshold for the topic publisher.<br>Index Type: PerMqTopicPublisher                    | 1600 | 2000 | 30 | FALSE |
| <b>MqTopicPublisherPublishRateLow</b><br>The publish rate is below the threshold for the topic publisher.<br>Index Type: PerMqTopicPublisher                    | 15   | 5    | 30 | FALSE |
| <b>MqTopicPublisherTotalsExpired</b><br>The topic publisher totals has expired.<br>Index Type: PerMqTopicPublisherTotals                                        | NaN  | NaN  | 30 | FALSE |
| <b>MqTopicPublisherTotalsPublishRateHigh</b><br>The publish rate exceeds the threshold for the topic publisher totals.<br>Index Type: PerMqTopicPublisherTotals | 1600 | 2000 | 30 | FALSE |
| <b>MqTopicPublisherTotalsPublishRateLow</b><br>The publish rate is below the threshold for the topic publisher totals.<br>Index Type: PerMqTopicPublisherTotals | 15   | 5    | 30 | FALSE |

|                                                                                                                                                                    |      |      |    |       |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|------|------|----|-------|
| <b>MqTopicSubscriberExpired</b><br>The topic subscriber has expired.<br>Index Type: PerMqTopicSubscriber                                                           | NaN  | NaN  | 30 | FALSE |
| <b>MqTopicSubscriberMessageRateHigh</b><br>The message rate exceeds the threshold for the topic subscriber.<br>Index Type: PerMqTopicSubscriber                    | 1600 | 2000 | 30 | FALSE |
| <b>MqTopicSubscriberMessageRateLow</b><br>The message rate is below the threshold for the topic subscriber.<br>Index Type: PerMqTopicSubscriber                    | 15   | 5    | 30 | FALSE |
| <b>MqTopicSubscriberTotalsExpired</b><br>The topic subscriber totals has expired.<br>Index Type: PerMqTopicSubscriberTotals                                        | NaN  | NaN  | 30 | FALSE |
| <b>MqTopicSubscriberTotalsMessageRateHigh</b><br>The message rate exceeds the threshold for the topic subscriber totals.<br>Index Type: PerMqTopicSubscriberTotals | 1600 | 2000 | 30 | FALSE |
| <b>MqTopicSubscriberTotalsMessageRateLow</b><br>The message rate is below the threshold for the topic subscriber totals.<br>Index Type: PerMqTopicSubscriberTotals | 15   | 5    | 30 | FALSE |
| <b>MqTopicSubscriptionCountHigh</b><br>The subscription count exceeds the threshold for the topic.<br>Index Type: PerMqTopic                                       | 1600 | 2000 | 30 | FALSE |
| <b>MqTopicSubscriptionCountLow</b><br>The subscription count is below the threshold for the topic.<br>Index Type: PerMqTopic                                       | 15   | 5    | 30 | FALSE |

---

## 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><br>The number of current deadlocks has exceeded its threshold.<br><b>Index Type(s):</b> PerServer<br><b>Metric:</b> DeltaNumber of Deadlocks       | 1           | 2           | 0        | TRUE    |
| <b>MssqlInstanceLatchWaitsHigh</b><br>The number of current latch waits has exceeded its threshold.<br><b>Index Type(s):</b> PerServer<br><b>Metric:</b> DeltaLatch Waits                | 15          | 30          | 0        | TRUE    |
| <b>MssqlInstanceLockWaitsHigh</b><br>The amount of seconds on lock waits has exceeded its threshold.<br><b>Index Type(s):</b> PerServer<br><b>Metric:</b> DeltaLock Waits                | 15          | 30          | 0        | TRUE    |
| <b>MssqlInstancePacketErrorsDetected</b><br>The amount of current packet errors has exceeded its threshold.<br><b>Index Type(s):</b> PerServer<br><b>Metric:</b> DeltaPacket Errors      | 1           | 2           | 0        | TRUE    |
| <b>MssqlInstanceSqlCpuUsedHigh</b><br>The percentage of CPU utilization on SQL processing has exceeded its threshold.<br><b>Index Type(s):</b> PerServer<br><b>Metric:</b> CPU Util      | 10          | 15          | 0        | TRUE    |
| <b>MssqlInstanceUsedMemoryHigh</b><br>The percentage of memory used by the SQL Server has exceeded its threshold.<br><b>Index Type(s):</b> PerServer<br><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><br>A collection was not able to be contacted for longer than the normal expiration window.<br><b>Index Type(s):</b> PerCollection<br><b>Metric:</b> Expired           | NaN         | NaN         | 30       | FALSE   |
| <b>MongoCollectionNumObjectsHigh</b><br>The number of objects for the collection exceeds a given threshold.<br><b>Index Type(s):</b> PerCollection<br><b>Metric:</b> numberOfObjects                | 1600        | 2000        | 30       | FALSE   |
| <b>MongoDatabaseDataSizeHigh</b><br>The database size for the database exceeds a given threshold.<br><b>Index Type(s):</b> PerDatabase<br><b>Metric:</b> dataSize                                   | 80000       | 100000      | 30       | FALSE   |
| <b>MongDatabaseExpired</b><br>The database was not able to be contacted for longer than the normal expiration window.<br><b>Index Type(s):</b> PerDatabase<br><b>Metric:</b> Expired                | NaN         | NaN         | 30       | FALSE   |
| <b>MongoInstanceExpired</b><br>The instance was not able to be contacted for longer than the normal expiration window.<br><b>Index Type(s):</b> PerInstance<br><b>Metric:</b> Expired               | 60          | 80          | 30       | FALSE   |
| <b>MongoInstanceNotConnected</b><br>The instance was not able to be contacted for longer than the normal expiration window.<br><b>Index Type(s):</b> PerInstance<br><b>Metric:</b> connectionStatus | NaN         | NaN         | 30       | FALSE   |
| <b>MongoInstanceOpenCursorsHigh</b><br>The number of Open Cursors for the Instance exceeds a given threshold.<br><b>Index Type(s):</b> PerInstance<br><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 |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|-------------|----------|---------|
| <b>MysqlBytesReceivedHigh</b><br>Executes a single warning and a single alarm if the amount of kilobytes received exceeds the specified threshold.<br>Index Type: PerServer<br>Metric: Received                                                                                                                     | 5             | 10          | na       | FALSE   |
| <b>MysqlBytesSentHigh</b><br>Executes a single warning and a single alarm if the amount of kilobytes sent exceeds the specified threshold.<br>Index Type: PerServer<br>Metric: Sent                                                                                                                                 | 5             | 10          | na       | FALSE   |
| <b>MysqlDelayedWritesHigh</b><br>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.<br>Index Type: PerServer<br>Metric: Delayed Writes | 1             | 2           | na       | FALSE   |
| <b>MysqlLocksWaitedHigh</b><br>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.<br>Index Type: PerServer<br>Metric: Table_locks_waited                                                      | 1             | 2           | na       | FALSE   |
| <b>MysqlQcacheLowMemPrunesHigh</b><br>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.<br>Index Type: PerServer<br>Metric: Qcache_lowmem_prunes                                                            | 1             | 2           | na       | FALSE   |

|                                                                                                                                                                                                                                                                                 |   |   |    |       |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---|---|----|-------|
| <b>MysqlSlowQueriesHigh</b><br>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.<br>Index Type: PerServer<br>Metric: Slow Queries                   | 1 | 2 | na | FALSE |
| <b>MysqlSlowThreadsHigh</b><br>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.<br>Index Type: PerServer<br>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><br>A master node's CPU usage is above the defined thresholds.<br><b>Index Type(s):</b> PerConnection<br><b>Metric:</b> Node Master - CPU %                            | 30          | 50          | 30       | FALSE   |
| <b>NodeMasterExpired</b><br>A master node has expired.<br><b>Index Type(s):</b> PerConnection<br><b>Metric:</b> Node Master - Expired                                                               | NaN         | NaN         | 30       | FALSE   |
| <b>NodeMasterRequestRateHigh</b><br>The request rate of a master node is above the defined thresholds.<br><b>Index Type(s):</b> PerConnection<br><b>Metric:</b> Node Requests - Requests Per Second | 1600        | 2000        | 30       | FALSE   |
| <b>NodeMasterResponseTimeHigh</b><br>The response time of a URL is above the defined thresholds.<br><b>Index Type(s):</b> PerConnection<br><b>Metric:</b> Node Requests - Avg Response Time         | 5           | 10          | 30       | FALSE   |

|                                                                                                                                                                                         |     |     |    |       |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----|-----|----|-------|
| <b>NodeProcessCpuUsageHigh</b><br>A worker node's CPU usage is above the defined thresholds.<br><b>Index Type(s):</b> PerConnection<br><b>Metric:</b> Node Processes - CPU Used %       | 5   | 50  | 30 | TRUE  |
| <b>NodeProcessExpired</b><br>A worker node has expired.<br><b>Index Type(s):</b> PerConnection<br><b>Metric:</b> Node Processes - Expired                                               | NaN | NaN | 30 | FALSE |
| <b>NodeProcessMemUsageHigh</b><br>A master node's memory usage has exceeded the defined limits.<br><b>Index Type(s):</b> PerConnection<br><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><br>Executes a single alert if the SQL database connection state is false.<br>Index Type: Per Database<br>Metric: Connected                                                                                    | NaN           | NaN         | 30       | FALSE   |
| <b>OraDatabaseQueryError</b><br>Executes a single alert if the if the last query state is false (an error).<br>Index Type: Per Database<br>Metric: Last Query Status                                                                           | NaN           | NaN         | 30       | FALSE   |
| <b>OraDatabaseResponseTimeHigh</b><br>Executes a single warning and a single alarm if the time (in milliseconds) to execute a SQL query exceeds the specified threshold.<br>Index Type: Per Database<br>Metric: ResponseTimeMilliSec           | 200           | 220         | 30       | FALSE   |
| <b>OraDatabaseSpaceUsedHigh</b><br>Executes a single warning and a single alarm if the percent utilization of the space allocated to the database exceeds the specified threshold.<br>Index Type: Per Database<br>Metric: PercentUsedSpace     | 80            | 90          | 30       | FALSE   |
| <b>OraDatabaseTablespaceUsedHigh</b><br>Executes a single warning and a single alarm if the percent utilization of the database used by the tablespace exceeds the specified threshold.<br>Index Type: Per Table Space<br>Metric: USED_PERCENT | 80            | 90          | 30       | FALSE   |
| <b>OraInstanceAvgQueryTimeHigh</b><br>Executes a single warning and a single alarm if the average time (in milliseconds) to perform a query exceeds the specified threshold.<br>Index Type: Per Instance<br>Metric: AVGQUERYTIME               | 300           | 400         | 30       | FALSE   |
| <b>OraInstanceCommitRateHigh</b><br>Executes a single warning and a single alarm if the number of commits per second exceeds the specified threshold.<br>Index Type: Per Instance<br>Metric: RateCOMMITTS                                      | 250           | 300         | 30       | FALSE   |

---

|                                                                                                                                                                                                                                    |       |       |    |       |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------|-------|----|-------|
| <b>OralInstanceNumCurrentLoginsHigh</b><br>Executes a single warning and a single alarm if the number of database clients exceeds the specified threshold.<br>Index Type: Per Instance<br>Metric: CURRENT_LOGINS                   | 12    | 15    | 30 | FALSE |
| <b>OralInstanceDataDictHitRatioLow</b><br>Executes a single warning and a single alarm if the data dictionary hit ratio goes below the specified threshold.<br>Index Type: Per Instance<br>Metric: DD_HIT_RATIO                    | 95    | 90    | 30 | FALSE |
| <b>OralInstanceDiskReadRateHigh</b><br>Executes a single warning and a single alarm if the number of physical disk reads per second exceeds the specified threshold.<br>Index Type: Per Instance<br>Metric: RatePHYSICAL_READS     | 250   | 300   | 30 | FALSE |
| <b>OralInstanceDiskWriteRateHigh</b><br>Executes a single warning and a single alarm if the number of physical disk writes per second exceeds the specified threshold.<br>Index Type: Per Instance<br>Metric: RatePHYSICAL_WRITES  | 250   | 300   | 30 | FALSE |
| <b>OralInstanceLatchHitRatioLow</b><br>Executes a single warning and a single alarm if the latch hit ratio goes below the specified threshold.<br>Index Type: Per Instance<br>Metric: LatchHitPerCent                              | 95    | 90    | 30 | FALSE |
| <b>OralInstanceMaxQueryTimeHigh</b><br>Executes a single warning and a single alarm if the query time (in milliseconds) exceeds the specified threshold.<br>Index Type: Per Instance<br>Metric: MAXQUERYTIME                       | 10000 | 15000 | 30 | FALSE |
| <b>OralInstanceNumActiveSessionsHigh</b><br>Executes a single warning and a single alarm if the number of active sessions for the instance exceeds the specified threshold.<br>Index Type: Per Instance<br>Metric: ACTIVE_SESSIONS | 12    | 15    | 30 | FALSE |
| <b>OralInstanceNumCurrentLoginsHigh</b><br>Executes a single warning and a single alarm if the number of current logins for the instance exceeds the specified threshold.<br>Index Type: Per Instance<br>Metric: CURRENT_LOGINS    | 12    | 15    | 30 | FALSE |

|                                                                                                                                                                                                                   |     |     |    |       |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----|-----|----|-------|
| <b>OracleInstanceRollbackRateHigh</b><br>Executes a single warning and a single alarm if the number of rollbacks per second exceeds the specified threshold.<br>Index Type: Per Instance<br>Metric: RateROLLBACKS | 5   | 10  | 30 | FALSE |
| <b>OracleInstanceSqlHitRatioLow</b><br>Executes a single warning and a single alarm if the SQL hit ratio goes below the specified threshold.<br>Index Type: Per Instance<br>Metric: SQL_HIT_RATIO                 | 95  | 90  | 30 | FALSE |
| <b>OracleInstanceState</b><br>Executes a single warning and a single alarm if the database is not in an ACTIVE or OPEN state for queries.<br>Index Type: Per Instance<br>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><br>The rate per second of newly opened sessions is below the specified threshold.<br><b>Index Type(s):</b> PerApplication                                                 | 10          | 1           | 30       | FALSE   |
| <b>WlsAppOpenSessionsHigh</b><br>The maximum total number of open sessions for that application has been reached.<br><b>Index Type(s):</b> PerApplication                                                 | 7           | 10          | 30       | FALSE   |
| <b>WlsClusterServersPercentNotRunningHigh</b><br>The percentage of cluster not running is high.<br><b>Index Type(s):</b> PerCluster                                                                       | 33          | 50          | 30       | FALSE   |
| <b>WlsHoggingThreadsHigh</b><br>The maximum number of hogging threads for that server has been reached.<br><b>Index Type(s):</b> PerServer                                                                | 15          | 20          | 30       | FALSE   |
| <b>WlsJDBCConnectionsWaitingHigh</b><br>Triggered when the number of threads waiting for a JDBC connection exceeds the threshold.<br><b>Index Type(s):</b> PerConnection, PerLocation, PerModule, PerName | 1           | 10          | 0        | FALSE   |

|                                                                                                                                                                                                 |    |    |    |       |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----|----|----|-------|
| <b>WlsJmsBytesCurrentHigh</b><br>The current number of bytes stored on this JMS server has reached its maximum.<br><b>Index Type(s):</b> PerServer                                              | 85 | 95 | 30 | FALSE |
| <b>WlsJmsBytesPendingHigh</b><br>The current number of bytes pending (unacknowledged or uncommitted) stored on this JMS Server has reached its maximum.<br><b>Index Type(s):</b> PerServer      | 85 | 95 | 30 | FALSE |
| <b>WlsJmsConnectionsCurrentHigh</b><br>The current number of connections to this JMS WebLogic Server has reached its maximum.<br><b>Index Type(s):</b> PerServer                                | 85 | 95 | 30 | FALSE |
| <b>WlsJmsDestinationBytesCurrentHigh</b><br>The current number of bytes stored in the destination, not including the pending bytes, has reached its maximum.<br><b>Index Type(s):</b> PerServer | 85 | 95 | 30 | FALSE |
| <b>WlsJmsDestinationBytesPendingHigh</b><br>The number of pending bytes stored in the destination has reached its maximum.<br><b>Index Type(s):</b> PerServer                                   | 85 | 95 | 30 | FALSE |
| <b>WlsJmsDestinationConsumersCurrentLow</b><br>The number of pending bytes stored in the destination has reached its minimum.<br><b>Index Type(s):</b> PerServer                                | 15 | 5  | 30 | FALSE |
| <b>WlsJmsDestinationMessagesCurrentHigh</b><br>The current number of messages in the destination has reached its maximum.<br><b>Index Type(s):</b> PerServer                                    | 85 | 95 | 30 | FALSE |
| <b>WlsJmsDestinationMessagesPendingHigh</b><br>The number of pending messages in the destination has reached its maximum.<br><b>Index Type(s):</b> PerServer                                    | 85 | 95 | 30 | FALSE |

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |     |     |    |       |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----|-----|----|-------|
| <b>WlsJmsDestinationsCurrentLow</b><br>The current number of destinations on this JMS Server has reached its minimum.<br><b>Index Type(s):</b> PerServer<br><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:<br><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 |
| <b>WlsJmsMessagesPendingHigh</b><br>The current number of messages pending (unacknowledged or uncommitted) stored on this JMS Server has reached its maximum.<br><b>Index Type(s):</b> PerServer                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | 85  | 95  | 30 | FALSE |
| <b>WlsJmsServerHealthNotOK</b><br>The health state of this JMS Server is not OK.<br><b>Index Type(s):</b> PerServer                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | NaN | NaN | 30 | FALSE |
| <b>WlsLockedUserCurrentHigh</b><br>The maximum number of current locked users for that server has been reached.<br><b>Index Type(s):</b> PerServer<br><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:<br><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 |
| <b>WlsOpenSocketsHigh</b><br>The maximum number of open sockets for that server has been reached.<br><b>Index Type(s):</b> PerServer                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | 85  | 95  | 30 | FALSE |
| <b>WlsPendingRequestCurrentHigh</b><br>The maximum number of current requests for that server has been reached.<br><b>Index Type(s):</b> PerServer                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | 85  | 95  | 30 | FALSE |

|                                                                                                                                                                                                                |     |     |    |       |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----|-----|----|-------|
| <b>WlsQueueLengthHigh</b><br>The number of pending requests in the priority queue has reached its maximum. This is the total of internal system requests and user requests.<br><b>Index Type(s):</b> PerServer | 85  | 95  | 30 | FALSE |
| <b>WlsServerCpuHigh</b><br>The server CPU has reached its maximum.<br><b>Index Type(s):</b> PerServer                                                                                                          | 85  | 95  | 30 | FALSE |
| <b>WlsServerHealthNotOK</b><br>The server health is not OK.<br><b>Index Type(s):</b> PerServer                                                                                                                 | NaN | NaN | 30 | FALSE |
| <b>WlsServerHostCpuHigh</b><br>The CPU percentage of the host server has reached its maximum.<br><b>Index Type(s):</b> PerServer                                                                               | 85  | 95  | 30 | FALSE |
| <b>WlsServerMemoryUsageHigh</b><br>The maximum used memory established for the server has been reached.<br><b>Index Type(s):</b> PerServer                                                                     | 85  | 95  | 30 | FALSE |
| <b>WlsServerNewSessionsLow</b><br>The number of new sessions created is below the threshold.<br><b>Index Type(s):</b> PerServer                                                                                | 15  | 5   | 30 | FALSE |
| <b>WlsServerOpenSessionsHigh</b><br>The maximum number of open sessions for that server has been reached.<br><b>Index Type(s):</b> PerServer                                                                   | 85  | 95  | 30 | FALSE |
| <b>WlsServerPendingUserRequestsHigh</b><br>The maximum number of pending user requests has been reached.<br><b>Index Type(s):</b> PerServer                                                                    | 85  | 95  | 30 | FALSE |
| <b>WlsServerReloadsHigh</b><br>The maximum number of reloads for that server has been reached.<br><b>Index Type(s):</b> PerServer                                                                              | 85  | 95  | 30 | FALSE |
| <b>WlsServerStaleData</b><br>The server has stale data.<br><b>Index Type(s):</b> PerServer                                                                                                                     | NaN | NaN | 30 | FALSE |
| <b>WlsServerStateNotRunning</b><br>The state of the server is different from "Running."<br><b>Index Type(s):</b> PerServer                                                                                     | NaN | NaN | 30 | FALSE |



|                                                                                                                                               |    |    |    |       |
|-----------------------------------------------------------------------------------------------------------------------------------------------|----|----|----|-------|
| <b>WlsThreadsTotalHigh</b><br>The total number of threads for that server has been reached.<br><b>Index Type(s):</b> PerServer                | 50 | 95 | 30 | FALSE |
| <b>WlsTransactionRolledBackTotalHigh</b><br>The total number of transactions rolled back has been reached.<br><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><br>Executes a single warning alert and a single alarm alert if the average CPU load per minute exceeds the specified threshold.<br><b>Index Type(s):</b> PerHost<br><b>Metric:</b> loadAvg1       | 50            | 75          | 30       | FALSE   |
| <b>HostCpuLoadAvg5High</b><br>Executes a single warning alert and a single alarm alert if the average CPU load per 5 minutes exceeds the specified threshold.<br><b>Index Type(s):</b> PerHost<br><b>Metric:</b> loadAvg5    | 50            | 75          | 30       | FALSE   |
| <b>HostCpuLoadAvg15High</b><br>Executes a single warning alert and a single alarm alert if the average CPU load per 15 minutes exceeds the specified threshold.<br><b>Index Type(s):</b> PerHost<br><b>Metric:</b> loadAvg15 | 50            | 75          | 30       | FALSE   |
| <b>HostCpuPercentHigh</b><br>Executes a single warning alert and a single alarm alert if the percent CPU load exceeds the specified threshold.<br><b>Index Type(s):</b> PerHost<br><b>Metric:</b> hostCpuPercent             | 50            | 75          | 30       | FALSE   |

|                                                                                                                                                                                                                                                 |     |     |    |       |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----|-----|----|-------|
| <b>HostMemoryUsedHigh</b><br>Executes a single warning alert and a single alarm alert if the percent of physical memory used exceeds the specified threshold.<br>Index Type(s): PerHost<br>Metric: MemUsedPerCent                               | 75  | 90  | 5  | FALSE |
| <b>HostNetworkRxRateHigh</b><br>Executes a single warning alert and a single alarm alert if the inbound network data rate, in kilobytes per second, exceeds the specified threshold.<br>Index Type(s): PerHost<br>Metric: RateRxKBytes          | 50  | 75  | 30 | FALSE |
| <b>HostNetworkTxRateHigh</b><br>Executes a single warning alert and a single alarm alert if the outbound network transmission rate, in kilobytes per second, exceeds the specified threshold.<br>Index Type(s): PerHost<br>Metric: RateTxKBytes | 50  | 75  | 30 | FALSE |
| <b>HostProcessCountLow</b><br>Executes a single warning alert and a single alarm alert if the process count exceeds the specified threshold.<br>Index Type(s): PerHost<br>Metric: Count                                                         | 80  | 90  | 30 | FALSE |
| <b>HostStaleData</b><br>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.<br>Index Type(s): PerHost<br>Metric: Expired          | NaN | NaN | 30 | FALSE |
| <b>HostStorageUsedHigh</b><br>Executes a single warning alert and a single alarm alert if the percent of space used on the storage medium exceeds the specified threshold.<br>Index Type(s): PerStorage<br>Metric: percentused                  | 80  | 90  | 5  | FALSE |
| <b>HostSwapUsedHigh</b><br>Executes a single warning alert and a single alarm alert if the percent of used swap space exceeds the specified threshold.<br>Index Type(s): PerHost<br>Metric: swapUsedPerCent                                     | 75  | 90  | 30 | FALSE |
| <b>HostVirtualMemoryUsedHigh</b><br>Executes a single warning alert and a single alarm alert if the percent of used virtual memory exceeds the specified threshold.<br>Index Type(s): PerHost<br>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.<br>Index Type: Per JVM<br>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.<br>Index Type: Per GC Source<br>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.<br>Index Type: Per GC Source<br>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.<br>Index Type(s): Per JVM<br>Metric: MemoryUsedPercent                                |
| <b>JvmNotConnected</b>          | Executes a single alert if the JVM is disconnected, indicating that it might have crashed.<br>Index Type(s): Per JVM<br>Metric: Connected                                                                     |
| <b>JvmStaleData</b>             | Executes a single alert if the data update wait time exceeds the specified duration threshold.<br>Index Type(s): Per JVM<br>Metric: Expired                                                                   |
| <b>JvmThreadCountHigh</b>       | Executes a single warning alert and a single alarm alert if the number of threads exceeds the specified threshold.<br>Index Type(s): Per JVM<br>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.<br>Index Type(s): Per Server<br>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.<br>Index Type(s): Per Server<br>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.<br>Index Type(s): Per Application<br>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.<br>Index Type(s): Per Application<br>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><br>The number of inbound bytes per second across the bridge has reached its maximum.<br>Index Type: PerBridge              | 8000000       | 10000000    | 30       | FALSE   |
| <b>SolBridgeInboundMsgRateHigh</b><br>The number of inbound messages per second across the bridge as a whole has reached its maximum.<br>Index Type: PerBridge | 40000         | 50000       | 30       | FALSE   |
| <b>SolBridgeOutboundByteRateHigh</b><br>The number of outbound bytes per second across the bridge has reached its maximum.<br>Index Type: PerBridge            | 8000000       | 10000000    | 30       | FALSE   |
| <b>SolBridgeOutboundMsgRateHigh</b><br>The number of outbound messages per second across the bridge has reached its maximum.<br>Index Type: PerBridge          | 40000         | 50000       | 30       | FALSE   |
| <b>SolBrokerEgressByteRateHigh</b><br>The egress rate (bytes/sec) for the message broker is excessive.<br>Index Type: PerAppliance                             | 70            | 85          | 30       | FALSE   |

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |    |    |    |       |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----|----|----|-------|
| <p><b>SolBrokerIngressByteRateHigh</b><br/>The ingress rate (bytes/sec) for the message broker is excessive.<br/>Index Type: PerAppliance</p>                                                                                                                                                                                                                                                                                                                               | 70 | 85 | 30 | FALSE |
| <p><b>SolBrokerEgressMsgRateHigh</b><br/>The message egress rate (the number of egress messages per second) for the message broker is excessive.<br/>Index Type: PerAppliance</p>                                                                                                                                                                                                                                                                                           | 70 | 85 | 30 | FALSE |
| <p><b>SolBrokerIngressMsgRateHigh</b><br/>The message ingress rate (the number of ingress messages per second) for the message broker is excessive.<br/>Index Type: PerAppliance</p>                                                                                                                                                                                                                                                                                        | 70 | 85 | 30 | FALSE |
| <p><b>SolBrokerNoQueueFound</b><br/>This is an Event Alert. Event Alerts do not have duration or threshold settings.<br/>A single alarm alert when there are discarded queues in the broker. (Delta of discard-queue-not-found is non-zero).<br/><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.<br/>Index Type: PerBroker</p>      |    |    |    | FALSE |
| <p><b>SolBrokerNoSubscriptionMatch</b><br/>This is an Event Alert. Event Alerts do not have duration or threshold settings.<br/>A single alarm alert when there are no current subscription matches (Delta of no-subscription-match is non-zero).<br/><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.<br/>Index Type: PerBroker</p> |    |    |    | FALSE |
| <p><b>SolBrokerNoValidDestination</b> This is an Event Alert. Event Alerts do not have duration or threshold settings.<br/>A single alarm alert when invalid destinations exist in the broker. (Delta of discard-nodest is non-zero).<br/><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.<br/>Index Type: PerBroker</p>             |    |    |    | FALSE |

**SolBrokerRedundancyDown & SolBrokerRedundancyActivityStatusChanged**

These alerts only pertain to brokers that are configured for redundancy.

These alerts execute when a redundancy misconfiguration is detected. Brokers qualify as being configured for redundancy if the Monitor either detects an associated mate broker name or a broker is explicitly configured for redundancy.

To verify whether the Monitor has detected all brokers configured for redundancy, go to the **Admin>Cache Table** display and select the **\_SolBrokerRedundancy** cache. Verify that all brokers that are configured for redundancy have the **ISHABroker** flag checked. If the **ISHABroker** flag is NOT checked, use the RTView Configuration Application to configure the brokers for redundancy.

The **SolBrokerRedundancyDown** alert verifies that redundancy is configured properly by checking whether **Redundancy Mode**, **Redundancy Status** and **Configuration Status** are valid. That is, the **Redundancy Mode** is either **Active/Active** or **Active/Standby**, the **Redundancy Status** is **Up** and the **Redundancy Configuration Status** is **Enabled**. If any of these conditions are not met, then a warning alert will be raised with the following alert text: "<hostname> is not properly configured for redundancy or redundancy is down. Redundancy Status: <a> Configuration Status: <b>", where <hostname> is the hostname of the offending broker and <a> and <b> are the current Redundancy Status and Configuration Status of the broker respectively.

The **SolBrokerRedundancyActivityStatusChanged** alert checks whether the previous state of the **Active-Standby Role**, the **Activity Status** of the Primary Virtual Router and the **Activity Status** of the Backup Virtual Router is different from the current state. If they are different, that implies a change in the state of the redundancy status occurred and a warning alert will be triggered. As soon as the previous and the current redundancy state is stabilized, the warning alert automatically clears, indicating in the alert text the current and previous states being detected. The warning alert contains the following text: "<hostname> has changed its redundancy activity state. There might be untracked intermediate states from the ones that have been detected. Current state: <A> Previous state: <B>", where <A> and <B> are the concatenation of active-standby-role, primary-status-activity, and backup-status-activity separated by the character "-" for current and previous states.

**Best Practices & Troubleshooting**

It's possible to have multiple **SolBrokerRedundancyActivityStatusChanged** warning alerts when failing over if intermediate states have been collected. For instance, if the changes from Local Active to Local Inactive to Shutdown are detected, then two **SolBrokerRedundancyActivityStatusChanged** warning alerts will be executed in this broker and will have two warnings from one broker and one from the other broker if the intermediate state on the second broker was not gathered due to polling interval being longer than the time the broker changes its redundancy state. If you only want one warning alert per broker per failover operation, the recommended action is to increase the duration of the alert. This value will vary depending on data collection latency and is system dependent. On the other hand, if you need to keep track of all intermediate states of the failover operation, then you should decrease the polling interval for the show redundancy detail poller. This is not recommended as might overflow the data collector with requests that cannot be successfully completed or preventing sending other monitoring data regarding other aspects of the broker due to the existence of requests too-often repeated.

Due to **SolBrokerRedundancyActivityStatusChanged** warning alert being a transient alert which will be automatically cleared when the redundancy status is stabilized, enabling both alerts is recommended as **SolBrokerRedundancyDown** can stay uncleared if manual intervention for fixing redundancy misconfiguration or non-functioning is required.

By default, these alerts are disabled.

|                                                                                                                                                             |         |          |    |       |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|----------|----|-------|
| <b>SolClientInboundByteRateHigh</b><br>The number of inbound bytes per second for the client has reached its maximum.<br>Index Type: PerClient              | 8000000 | 10000000 | 30 | FALSE |
| <b>SolClientInboundMsgRateHigh</b><br>The number of inbound messages per second for the client as a whole has reached its maximum.<br>Index Type: PerClient | 40000   | 50000    | 30 | FALSE |
| <b>SolClientOutboundByteRateHigh</b><br>The number of outbound bytes per second for the client has reached its maximum.<br>Index Type: PerClient            | 8000000 | 10000000 | 30 | FALSE |

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |       |       |     |       |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------|-------|-----|-------|
| <b>SolClientOutboundMsgRateHigh</b><br>The number of outbound messages per second for the client as a whole has reached its maximum.<br>Index Type: PerClient                                                                                                                                                                                                                                                                                                                                                                                                                                                      | 40000 | 50000 | 30  | FALSE |
| <b>SolClientSlowSubscriber</b><br>One or more clients are consuming messages too slowly; endpoints may drop messages!<br>Index Type: PerClient                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | 1     | NaN   | 30  | FALSE |
| <b>SolCspfNeighborDown</b><br>State is not "OK" for one or more CSPF neighbors.<br>Index Type: PerNeighbor                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | 1     | NaN   | 30  | FALSE |
| <b>SolEndpointNoBridgeClient</b><br>This is an Event Alert. Event Alerts do not have duration or threshold settings.<br>A single alarm alert when there are no binds for the Solace Endpoint exist (bind-count is zero).<br>Index Type: PerEndpoint                                                                                                                                                                                                                                                                                                                                                                | NaN   | NaN   | NaN | FALSE |
| <b>SolEndpointNoBridgeTopic</b><br>This is an Event Alert. Event Alerts do not have duration or threshold settings.<br>A single alarm alert when there are no topics subscribed to the Queue (topic-subscription-count is zero).<br>Index Type: PerEndpoint                                                                                                                                                                                                                                                                                                                                                        | NaN   | NaN   | NaN | FALSE |
| <b>SolEndpointPendingMsgsHigh</b><br>The number of pending messages on a queue has reached its maximum.<br>Index Type: PerEndpoint                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | 8000  | 10000 | 30  | FALSE |
| <b>SolEndpointSpoolUsageHigh</b><br>The endpoint is consuming too much message broker memory for storing spooled messages. (Threshold units are megabytes.)<br>Index Type: PerEndpoint                                                                                                                                                                                                                                                                                                                                                                                                                             | 40    | 50    | 30  | FALSE |
| <b>SolEventModuleBrokerAlert</b><br>This is an Event Alert. Event Alerts do not have duration or threshold settings.<br>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><br>This is an Event Alert. Event Alerts do not have duration or threshold settings.<br>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><br>This is an Event Alert. Event Alerts do not have duration or threshold settings.<br>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><br>For Guaranteed Messaging only, the Operational State for each HBA Fibre-Channel should be Online (e.g., not Linkdown).<br>Index Type: PerHbaLink                                                                                                                                                                                                                                                                                                                                                                                                                 | NaN | 0   | 30  |  | FALSE |
| <b>SolGuaranteedMsgingMatePortDown</b><br>This alert executes when the broker is configured for Guaranteed Messaging and the ADB Link to Mate is down.<br>Note that this alert does not clear until a failover occurs.<br>Index Type: PerADB                                                                                                                                                                                                                                                                                                                                                              | NaN | 0   | 30  |  | FALSE |
| <b>SolGuaranteedMsgingNoMsgSpoolAdActive</b><br>This alert applies to a pair of brokers that are configured for redundancy as an HA pair. A single alert executes when neither broker in the HA pair has a message spool operational status of <b>AD-Active</b> .<br>Index Type: PerMsgRouter                                                                                                                                                                                                                                                                                                             | NaN | 0   | 30  |  | FALSE |
| <b>SolMsgBrokerExpired</b><br>This is an Event Alert. Event Alerts do not have duration or threshold settings.<br>The collection of monitoring data has stopped abruptly. Use this alert in conjunction with the <b>SolMsgBrokerNotConnected</b> alert, or instead of it, if you don't need to be notified about a lack of connection when the Monitor starts up.                                                                                                                                                                                                                                         | NaN | NaN | NaN |  | FALSE |
| <b>SolMsgRouterActiveDiskUtilHigh</b><br>The utilization of the active disk partition for the message broker is excessive.<br>Index Type: PerAppliance                                                                                                                                                                                                                                                                                                                                                                                                                                                    | 70  | 85  | 30  |  | FALSE |
| <b>SolMsgRouterConnectionUtilHigh</b><br>The connection utilization for the message broker (current number of connections divided by max allowed) is excessive.<br>Index Type: PerAppliance                                                                                                                                                                                                                                                                                                                                                                                                               | 70  | 85  | 30  |  | FALSE |
| <b>SolMsgRouterCpuTemperatureHigh</b><br>CPU temperature margin is above threshold.<br>Index Type: PerApplianceSensor                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | -30 | -15 | 30  |  | FALSE |
| <b>SolMsgRouterCspfNeighborDown</b><br>Link-detect = no for CSPF neighbor.<br>Index Type: PerAppliance                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | 1   | NaN | 30  |  | FALSE |
| <b>SolMsgRouterDelvrUnAckMsgUtilHigh</b><br>The delivered unacked messages as a percentage of all messages delivered for the application is excessive.<br>Index Type: PerAppliance                                                                                                                                                                                                                                                                                                                                                                                                                        | 70  | 85  | 30  |  | FALSE |



|                                                                                                                                                                                |        |        |    |       |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------|--------|----|-------|
| <b>SolMsgRouterFanSensorCheckFailed</b><br>The speed measured for one or more fans is below threshold.<br>Index Type: PerApplianceSensor                                       | 5000   | 2657   | 30 | FALSE |
| <b>SolMsgRouterInboundByteRateHigh</b><br>The number of inbound bytes per second for the message broker has reached its max threshold.<br>Index Type: PerAppliance             | 400000 | 500000 | 30 | FALSE |
| <b>SolMsgRouterInboundMsgRateHigh</b><br>The number of inbound messages per second for the message broker has reached its max threshold.<br>Index Type: PerAppliance           | 400000 | 500000 | 30 | FALSE |
| <b>SolMsgRouterIngressFlowUtilHigh</b><br>The ingress flow utilization (current flows divided by max allowed) for the message broker is excessive.<br>Index Type: PerAppliance | 70     | 85     | 30 | FALSE |
| <b>SolMsgRouterInterfaceDown</b><br>Link-detect = no for one or more enabled network interfaces.<br>Index Type: PerSolInterface                                                | NaN    | NaN    | 30 | FALSE |
| <b>SolMsgRouterMsgCountUtilHigh</b><br>The message count utilization for the message broker is excessive.<br>Index Type: PerAppliance                                          | 70     | 85     | 30 | FALSE |
| <b>SolMsgRouterNABUsageHigh</b><br>Network Acceleration Blade memory usage is excessive.<br>Index Type: PerNAB                                                                 | 60     | 80     | 30 | FALSE |
| <b>SolMsgRouterNotConnected</b><br>The message broker is not ready for collecting performance monitoring data.<br>Index Type: PerAppliance                                     | NaN    | NaN    | 30 | FALSE |
| <b>SolMsgRouterOutboundByteRateHigh</b><br>The number of outbound bytes per second for the message broker has reached its max threshold.<br>Index Type: PerAppliance           | 400000 | 500000 | 30 | FALSE |
| <b>SolMsgRouterOutboundMsgRateHigh</b><br>The number of outbound messages per second for the message broker has reached its max threshold.<br>Index Type: PerAppliance         | 400000 | 500000 | 30 | FALSE |
| <b>SolMsgRouterPendingMsgsHigh</b><br>The total number of pending messages for this message broker has reached its maximum.<br>Index Type: PerAppliance                        | 400000 | 500000 | 30 | FALSE |
| <b>SolMsgRouterPowerSupplyFailed</b><br>A power supply has failed.<br>Index Type: PerAppliance                                                                                 | 0      | NaN    | 30 | FALSE |

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |     |     |    |       |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----|-----|----|-------|
| <b>SolMsgRouterSpoolUtilization</b><br>The percentage of spool spaces used for storing spooled messages is excessive.<br>Index Type: PerAppliance                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | 70  | 85  | 30 | FALSE |
| <b>SolMsgRouterStandbyDiskUtilHigh</b><br>The utilization of the standby disk partition for the message broker is excessive.<br>Index Type: PerAppliance                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | 70  | 85  | 30 | FALSE |
| <b>SolMsgRouterSubscriptionUtilHigh</b><br>The Topic Subscriptions on Queue utilization (current number of subscriptions divided by max allowed) for the message router is excessive.<br>Index Type: PerAppliance                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | 70  | 85  | 30 | FALSE |
| <b>SolMsgRouterSwapUsedHigh</b><br>This alert strictly applies to software brokers (it does not execute for hardware brokers). The amount of swap space used by the message broker operating system is excessive.<br>Index Type: PerAppliance                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | 70  | 85  | 30 | FALSE |
| <b>SolMsgRouterSyslogAlert</b><br>This alert executes when a Solace Syslog Warning or Critical message is received. To get Syslog event alerts, go to the Alert Administration display and enable the <b>SolMsgRouterSyslog</b> alert.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | -   | -   | -  | -     |
| <b>SolMsgRouterTemperatureSensorCheckFailed</b><br>A chassis temperature measurement is above threshold.<br>Index Type: PerAppliance                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | 40  | 45  | 30 | FALSE |
| <b>SolMsgRouterTranSessionCntUtilHigh</b><br>The transacted session count utilization for the message broker is excessive. The metrics are:<br>(transacted-sessions-used/ max-transacted-sessions)*100<br>Index Type: PerMsgRouter                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | 70  | 85  | 30 | FALSE |
| <b>SolMsgRouterTranSessionResUtilHigh</b><br>The transacted session resource utilization for the message broker is excessive.<br>Index Type: PerAppliance                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | 70  | 85  | 30 | FALSE |
| <b>SolMsgRouterVoltageSensorCheckFailed</b><br>A power supply voltage is high or low.<br>Index Type: PerApplianceSesor                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | NaN | NaN | 30 | FALSE |
| <b>SolSparseMessageSpoolFile</b><br>This is a Limits Alert that issues a Warning alert and is enabled by default.<br><b>Important:</b> Do not modify thresholds for this alert as are set by Solace development.<br>A single warning alert (Severity 1) executes when the active-disk-partition-usage > 30.0<br>AND<br>disk-usage-mb/current-disk-usage >= 3.0.<br>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 for diagnosing the sparse message spool file condition. |     |     |    | TRUE  |

|                                                                                                                                                         |         |          |    |       |
|---------------------------------------------------------------------------------------------------------------------------------------------------------|---------|----------|----|-------|
| <b>SolVpnConnectionCountHigh</b><br>The number of connections to the server has reached its maximum.<br>Index Type: PerVPN                              | 60      | 80       | 30 | FALSE |
| <b>SolVpnInboundByteRateHigh</b><br>The number of inbound bytes per second for the VPN has reached its maximum.<br>Index Type: PerVPN                   | 8000000 | 10000000 | 30 | FALSE |
| <b>SolVpnInboundDiscardRateHigh</b><br>The number of discarded inbound messages per second for the server is excessive.<br>Index Type: PerVPN           | 1       | 5        | 30 | FALSE |
| <b>SolVpnInboundMsgRateHigh</b><br>The number of inbound messages per second for the VPN as a whole has reached its maximum.<br>Index Type: PerVPN      | 40000   | 50000    | 30 | FALSE |
| <b>SolVpnOutboundByteRateHigh</b><br>The number of outbound bytes per second for the VPN has reached its maximum.<br>Index Type: PerVPN                 | 8000000 | 10000000 | 30 | FALSE |
| <b>SolVpnOutboundDiscardRateHigh</b><br>The number of discarded outbound messages per second for the server is excessive.<br>Index Type: PerVPN         | 1       | 5        | 30 | FALSE |
| <b>SolVpnOutboundMsgRateHigh</b><br>The number of outbound messages per second for the server as a whole has reached its maximum.<br>Index Type: PerVPN | 40000   | 50000    | 30 | FALSE |
| <b>SolVpnPendingMsgsHigh</b><br>The total number of pending messages for this destination has reached its maximum.<br>Index Type: PerVPN                | 8000000 | 10000000 | 30 | FALSE |
| <b>SolVpnSubscriptionCountHigh</b><br>The number of endpoints in this VPN has reached its maximum.<br>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><br>BW6 application status is not Running or Stopped (status is Impaired, AppError or StartFailed)<br>Index Type: PerApp<br>Metric: State                                                                   | NaN           | NaN         | 30       | FALSE   |
| <b>Bw6AppExpired</b><br>BW6 application expired due to application inactivity.<br>Index Type: PerApp<br>Metric: Stopped                                                                                                            | NaN           | NaN         | 30       | FALSE   |
| <b>Bw6AppNodeCpuUsedHigh</b><br>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.<br>Index Type: PerAppNode<br>Metric: CPU Usage% | 50            | 80          | 30       | FALSE   |
| <b>Bw6AppNodeMemUsedHigh</b><br>BW6 AppNode memory usage exceeded limit. Memory usage is the percentage of total JVM memory currently consumed by this appnode.<br>Index Type: PerAppNode<br>Metric: Memory Usage%                 | 50            | 80          | 30       | FALSE   |
| <b>Bw6AppNodeStopped</b><br>BW6 AppNode stopped purposefully (for example, an administrator stopped the AppNode process).<br>Index Type: PerAppNode<br>Metric: State                                                               | NaN           | NaN         | 10       | FALSE   |
| <b>Bw6AppNodeUnreachable</b><br>BW6 AppNode stopped abnormally (for example, the AppNode process crashed).<br>Index Type: PerAppNode<br>Metric: State                                                                              | NaN           | NaN         | 10       | FALSE   |
| <b>Bw6AppProcessCreatedRateHigh</b><br>BW6 Process created rate for application exceeded limit.<br>Index Type: PerApp<br>Metric: App Created Rate                                                                                  | 50            | 80          | 30       | FALSE   |
| <b>Bw6AppProcessElapsedTimeHigh</b><br>BW6 Process delta elapsed time rate of increase for application exceeded limit.<br>Index Type: PerApp<br>Metric: App Elapsed Rate                                                           | 200           | 400         | 30       | FALSE   |

|                                                                                                                                                                                |     |     |    |       |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----|-----|----|-------|
| <b>Bw6AppProcessExecutionTimeHigh</b><br>BW6 Process delta execution time rate of increase for application exceeded limit.<br>Index Type: PerApp<br>Metric: App Execution Rate | 200 | 400 | 30 | FALSE |
| <b>Bw6AppProcessFailedRateHigh</b><br>BW6 Process failed rate for application exceeded limit.<br>Index Type: PerApp<br>Metric: App Failed Rate                                 | 50  | 80  | 30 | FALSE |
| <b>Bw6AppStopped</b><br>BW6 application stopped.<br>Index Type: PerApp<br>Metric: Stopped                                                                                      | NaN | NaN | 30 | FALSE |
| <b>Bw6ProcessActivityErrorRateHigh</b><br>BW6 Process error rate exceeded limit.<br>Index Type: PerProcess<br>Metric: Process Failed Rate                                      | 50  | 80  | 30 | FALSE |
| <b>Bw6ProcessCreatedRateHigh</b><br>BW6 Process error rate exceeded limit.<br>Index Type: PerProcess<br>Metric: Process Failed Rate                                            | 50  | 80  | 30 | FALSE |
| <b>Bw6ProcessElapsedTimeHigh</b><br>BW6 Process delta elapsed time rate of increase exceeded limit.<br>Index Type: PerProcess<br>Metric: Delta Exec Rate                       | 200 | 400 | 30 | FALSE |
| <b>Bw6ProcessExecutionTimeHigh</b><br>BW6 Process delta execution time rate of increase exceeded limit.<br>Index Type: PerProcess<br>Metric: Delta Time Rate                   | 200 | 400 | 30 | FALSE |
| <b>Bw6ProcessFailedRateHigh</b><br>BW6 Process suspended rate exceeded limit.<br>Index Type: PerProcess<br>Metric: Suspended Rate                                              | 50  | 80  | 30 | FALSE |
| <b>Bw6ProcessHung</b><br>The delta elapsed time is greater than zero but the delta execution time is zero.<br>Index Type: PerProcess<br>Metric: Hung/Not Hung                  | NaN | NaN | 10 | FALSE |
| <b>Bw6ProcessSuspendRateHigh</b><br>BW6 Process failed rate exceeded limit.<br>Index Type: PerProcess<br>Metric: Failed Rate                                                   | 50  | 80  | 30 | FALSE |

|                                                                                                                                                                                                                                                                                                                        |     |     |    |       |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----|-----|----|-------|
| <b>BwActivityErrorRateHigh</b><br>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.<br>Index Type: PerActivity<br>Metric: RateErrorCount                                                       | 50  | 80  | 30 | FALSE |
| <b>BwActivityExecutionTimeHigh</b><br>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.<br>Index Type: PerActivity<br>Metric: RateExecutionTime                           | 200 | 400 | 30 | FALSE |
| <b>BwEngineCpuUsedHigh</b><br>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.<br>Index Type: PerEngine<br>Metric: CPU Usage%                                                                               | 50  | 80  | 30 | FALSE |
| <b>BwEngineMemUsedHigh</b><br>BW Engine memory usage (% of total) exceeded limit. Memory usage is the percentage of total JVM memory currently consumed by this engine.<br>Index Type: PerEngine<br>Metric: PercentUsed                                                                                                | 50  | 80  | 30 | FALSE |
| <b>BwEngineStopped</b><br>BW Engine has stopped running.<br>Index Type: PerEngine<br>Metric: Stopped                                                                                                                                                                                                                   | NaN | NaN | 30 | FALSE |
| <b>BwEngineUnreachable</b><br>BW engine stopped abnormally.<br>Index Type: PerEngine<br>Metric: State                                                                                                                                                                                                                  | NaN | NaN | 30 | FALSE |
| <b>BwProcessAbortRateHigh</b><br>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.<br>Index Type: PerProcess<br>Metric: RateAborted                                                                   | 50  | 80  | 30 | FALSE |
| <b>BwProcessAvgElapsedTimeHigh</b><br>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.<br>Index Type: PerProcess<br>Metric: Process Avg Elapsed Time | 100 | 200 | 30 | FALSE |
| <b>BwProcessAvgExecutionTimeHigh</b><br>BW Process average execution time exceeded limit.<br>Index Type: PerProcess<br>Metric: AverageExecution                                                                                                                                                                        | 0   | 0   | 0  | FALSE |

|                                                                                                                                                                                                                                                                                            |     |     |    |       |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----|-----|----|-------|
| <b>BwProcessCreatedRateHigh</b><br>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.<br>Index Type: PerProcess<br>Metric: Processes Created/sec     | 100 | 200 | 30 | FALSE |
| <b>BwProcessCreatedRateLow</b><br>BW Process creation rate per second went below limit.<br>Index Type: PerProcess<br>Metric: App Created Rate                                                                                                                                              | 0   | 0   | 0  | FALSE |
| <b>BwProcessElapsedTimeHigh</b><br>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.<br>Index Type: PerProcess<br>Metric: RateTotalElapsed         | 50  | 80  | 30 | FALSE |
| <b>BwProcessExecutionTimeHigh</b><br>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.<br>Index Type: PerProcess<br>Metric: RateTotalExecution | 50  | 80  | 30 | FALSE |
| <b>BwProcessHung</b><br>The delta elapsed time is greater than zero but the delta execution time is zero.<br>Index Type: PerProcess<br>Metric: Hung/Not Hung                                                                                                                               | NaN | NaN | 10 | FALSE |
| <b>BwProcessSuspendRateHigh</b><br>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.<br>Index Type: PerProcess<br>Metric: RateSuspended                               | 50  | 80  | 30 | FALSE |
| <b>BwProcessTotalCpuPercentHigh</b><br>BW Process CPU percent utilization exceeded limit. This is the percent CPU used by all process instances executing over the interval.<br>Index Type: PerProcess<br>Metric: Process Total CPU Percent                                                | 50  | 75  | 30 | FALSE |
| <b>BwServerCpuUsedHigh</b><br>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.<br>Index Type: PerServer<br>Metric: CPU Usage%                                                 | 60  | 85  | 30 | FALSE |

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |     |     |    |       |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----|-----|----|-------|
| <p><b>BwServerFreeMemLow</b><br/>         BW Server free memory available (in megabytes) is below limit. Free memory means available physical (RAM) memory.<br/>         Index Type: PerServer<br/>         Metric: Memory Free Mbytes</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | 15  | 5   | 30 | FALSE |
| <p><b>BwServerInactive</b><br/>         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.<br/>         Index Type: PerServer<br/>         Metric: Expired</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | NaN | NaN | 30 | FALSE |
| <p><b>BwServerMemUsedHigh</b><br/>         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.<br/>         Index Type: PerServer<br/>         Metric: Virtual Memory Used%</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | 50  | 80  | 30 | FALSE |
| <p><b>HawkAlert</b><br/>         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.<br/>         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.<br/> <b>sl.rtvview.sub=\$hawkAlertTextFilterOut:AlertText</b><br/>         For example, to filter out all Hawk Alerts in which the AlertText contains <b>Source</b> you would enter the following:<br/> <b>sl.rtvview.sub=\$hawkAlertTextFilterOut:Source</b><br/>         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>:<br/> <b>sl.rtvview.sub=\$hawkAlertTextFilterOut:3600</b><br/>         Index Type: PerServer<br/>         Metric: Hawk</p> | NaN | NaN | -1 | TRUE  |
| <p><b>JvmCpuPercentHigh</b><br/>         The percentage of CPU that has been reached by the JVM is above the limit.<br/>         Index Type: PerJVM<br/>         Metric: CpuPercent</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | 50  | 75  | 30 | FALSE |
| <p><b>JvmGcDutyCycleHigh</b><br/>         The duty cycle is out the upper limit.<br/>         Index Type: PerGC<br/>         Metric: DutyCycle</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | 50  | 75  | 30 | FALSE |



|                                                                                                                    |     |     |    |       |
|--------------------------------------------------------------------------------------------------------------------|-----|-----|----|-------|
| <b>JvmMemoryUsedHigh</b><br>The memory used out the upper limit<br>Index Type: PerJVM<br>Metric: MemoryUsedPercent | 50  | 75  | 30 | FALSE |
| <b>JvmNotConnected</b><br>The JVM in not connected.<br>Index Type: PerJVM<br>Metric: Connected                     | NaN | NaN | 30 | FALSE |
| <b>JvmStaleData</b><br>Cut in reception from that JVM.<br>Index Type: PerJVM<br>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.<br>LEVEL | ALARM<br>LEVEL | DURATION | ENABLED |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------|----------------|----------|---------|
| <b>TdgKeeperCpuUsageHigh</b><br>The keeper CPU usage rate (msec/sec) is above the defined thresholds.<br><b>Index Type(s):</b> PerTdgKeeper                                | 60             | 80             | 30       | FALSE   |
| <b>TdgKeeperExpired</b><br>RTView is not receiving metrics updates from this Keeper. The Expired flag of the Keeper was set to true.<br><b>Index Type(s):</b> PerTdgKeeper | NaN            | NaN            | 30       | FALSE   |
| <b>TdgKeeperMemoryUseHigh</b><br>The keeper's usage of memory, in KB, is above the threshold.<br><b>Index Type(s):</b> PerTdgKeeper                                        | 1600000        | 2000000        | 30       | FALSE   |
| <b>TdgKeeperMsgsRcvdRateHigh</b><br>The incoming message rate, in messages per second, is higher than expected for this keeper.<br><b>Index Type(s):</b> PerTdgKeeper      | 160000         | 200000         | 30       | FALSE   |
| <b>TdgKeeperMsgsSentRateLow</b><br>The keeper's rate of messages sent is below the threshold.<br><b>Index Type(s):</b> PerTdgKeeper                                        | 15             | 5              | 30       | FALSE   |

|                                                                                                                                                                    |         |         |    |       |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|---------|----|-------|
| <b>TdgKeeperNotRunning</b><br>The current status for this keeper is not "RUNNING."<br><b>Index Type(s):</b> PerTdgKeeper                                           | NaN     | NaN     | 30 | FALSE |
| <b>TdgNodeCpuUsageHigh</b><br>The node CPU Usage rate (msec/sec) is above threshold.<br><b>Index Type(s):</b> PerTdgNode                                           | 60      | 80      | 30 | FALSE |
| <b>TdgNodeExpired</b><br>RTView is not receiving metrics updates from this Node. The Expired flag of the Node was set to true.<br><b>Index Type(s):</b> PerTdgNode | NaN     | NaN     | 30 | FALSE |
| <b>TdgNodeLiveDataSizeHigh</b><br>The node's live data size is above the threshold.<br><b>Index Type(s):</b> PerTdgNode                                            | 1600000 | 2000000 | 30 | FALSE |
| <b>TdgNodeMemoryUseHigh</b><br>The node's usage of memory, in KB, is above the threshold.<br><b>Index Type(s):</b> PerTdgNode                                      | 1600000 | 2000000 | 30 | FALSE |
| <b>TdgNodeMsgsRcvdRateHigh</b><br>The incoming message rate, in messages per second, is higher than expected for this node.<br><b>Index Type(s):</b> PerTdgNode    | 160000  | 200000  | 30 | FALSE |
| <b>TdgNodeMsgsSentRateLow</b><br>The outgoing message rate, in messages per second, is lower than expected for this node.<br><b>Index Type(s):</b> PerTdgNode      | 15      | 5       | 30 | FALSE |
| <b>TdgNodeNotRunning</b><br>The current status for this node is not "RUNNING".<br><b>Index Type(s):</b> PerTdgNode                                                 | NaN     | NaN     | 30 | FALSE |
| <b>TdgNodeOpsCompletedRateLow</b><br>The rate of completed operations on the node is below the threshold.<br><b>Index Type(s):</b> PerTdgNode                      | 15      | 5       | 30 | FALSE |
| <b>TdgNodeOpsFailedRateHigh</b><br>The rate of failed operations on the node is above the threshold.<br><b>Index Type(s):</b> PerTdgNode                           | 10      | 20      | 30 | FALSE |
| <b>TdgNodeTxnRollbackRateHigh</b><br>The node's rate of transactions rolled back is above the threshold.<br><b>Index Type(s):</b> PerTdgNode                       | 50      | 100     | 30 | FALSE |
| <b>TdgProxyCpuUsageHigh</b><br>The proxy CPU Usage rate (msec/sec) is above the defined threshold.<br><b>Index Type(s):</b> PerTdgProxy                            | 60      | 80      | 30 | FALSE |

|                                                                                                                                                                        |         |         |    |       |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|---------|----|-------|
| <b>TdgProxyExpired</b><br>RTView is not receiving metrics updates from this Proxy. The Expired flag of the Proxy was set to true.<br><b>Index Type(s):</b> PerTdgProxy | NaN     | NaN     | 30 | FALSE |
| <b>TdgProxyMemoryUseHigh</b><br>The proxy's usage of memory, in kilobytes, is above the threshold.<br><b>Index Type(s):</b> PerTdgProxy                                | 1600000 | 2000000 | 30 | FALSE |
| <b>TdgProxyMsgsRcvdRateHigh</b><br>The incoming message rate, in messages per second, is higher than expected for this proxy.<br><b>Index Type(s):</b> PerTdgProxy     | 160000  | 200000  | 30 | FALSE |
| <b>TdgProxyMsgsSentRateLow</b><br>The outgoing message rate, in messages per second, is lower than expected for this proxy.<br><b>Index Type(s):</b> PerTdgProxy       | 15      | 5       | 30 | FALSE |
| <b>TdgProxyNotRunning</b><br>The current status for this proxy is not "RUNNING."<br><b>Index Type(s):</b> PerTdgProxy                                                  | NaN     | NaN     | 30 | FALSE |
| <b>TdgProxyTxnRollbackRateHigh</b><br>The proxy's rate of transactions rolled back is above the threshold.<br><b>Index Type(s):</b> PerTdgProxy                        | 50      | 100     | 30 | FALSE |
| <b>TdgRealmOpsCompletedRateLow</b><br>The rate of completed operations on the realm is below the threshold.<br><b>Index Type(s):</b> PerTdgRealm                       | 15      | 5       | 30 | FALSE |
| <b>TdgRealmOpsFailedRateHigh</b><br>The rate of failed operations on the realm is above the threshold.<br><b>Index Type(s):</b> PerTdgRealm                            | 10      | 20      | 30 | FALSE |
| <b>TdgRealmServerCpuUsageHigh</b><br>The CPU utilization of the Realm Server, as a percentage, is above the threshold.<br><b>Index Type(s):</b> PerTdgRealm            | 60      | 80      | 30 | FALSE |
| <b>TdgRealmServerExpired</b><br>RTView is not receiving metrics updates from this Realm Server. The Expired flag was set to true.<br><b>Index Type(s):</b> PerTdgRealm | NaN     | NaN     | 30 | FALSE |
| <b>TdgRealmServerMemoryUseHigh</b><br>The Realm Server memory usage (RSS) is above threshold. Units are kilobytes.<br><b>Index Type(s):</b> PerTdgRealm                | 160     | 200     | 30 | FALSE |
| <b>TdgRealmTxnRollbackRateHigh</b><br>The node's rate of transactions rolled back is above the threshold.<br><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.<br>LEVEL | ALARM<br>LEVEL | DURATION | ENABLED |
|--------------------------------------------------------------------------------------------------------------------------------------------------|----------------|----------------|----------|---------|
| <b>TasMemberCpuHigh</b><br>The CPU usage is above the defined thresholds.<br><b>Index Type(s):</b> PerMember                                     | 80             | 95             | 30       | FALSE   |
| <b>TasMemberEntriesHigh</b><br>The number of objects inserted into the space is above the defined thresholds.<br><b>Index Type(s):</b> PerMember | 8000           | 10000          | 30       | FALSE   |
| <b>TasMemberEvictsRateHigh</b><br>The rate at which 'evicts' are occurring is above the defined thresholds.<br><b>Index Type(s):</b> PerMember   | 80             | 100            | 30       | FALSE   |
| <b>TasMemberExpireRateHigh</b><br>The rate at which 'expires' are occurring is above the defined thresholds.<br><b>Index Type(s):</b> PerMember  | 80             | 100            | 30       | FALSE   |
| <b>TasMemberGetRateHigh</b><br>The rate at which 'gets' are occurring is above the defined thresholds.<br><b>Index Type(s):</b> PerMember        | 80             | 100            | 30       | FALSE   |
| <b>TasMemberJvmMemoryUsedHigh</b><br>The percent JVM memory used is above the defined thresholds.<br><b>Index Type(s):</b> PerMember             | 80             | 95             | 30       | FALSE   |
| <b>TasMemberMemoryUsedHigh</b><br>The percent memory used is above the defined thresholds.<br><b>Index Type(s):</b> PerMember                    | 80             | 95             | 30       | FALSE   |
| <b>TasMemberPutRateHigh</b><br>The rate at which 'puts' are occurring is above the defined thresholds.<br><b>Index Type(s):</b> PerMember        | 80             | 100            | 30       | FALSE   |

|                                                                                                                                                                                                                                             |      |        |    |       |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------|--------|----|-------|
| <b>TasMemberSeederCapacity</b>                                                                                                                                                                                                              | 80   | 90     | 30 | FALSE |
| 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.<br><b>Index Type(s):</b> PerMemberandSpace |      |        |    |       |
| <b>TasMemberTakeRateHigh</b>                                                                                                                                                                                                                | 80   | 100    | 30 | FALSE |
| The rate at which 'takes' are occurring is above the defined thresholds.<br><b>Index Type(s):</b> PerMember                                                                                                                                 |      |        |    |       |
| <b>TasMetaspaceEntriesHigh</b>                                                                                                                                                                                                              | 8000 | 100000 | 30 | FALSE |
| The number of objects inserted into the metaspace is above the defined thresholds.<br><b>Index Type(s):</b> PerMetaspace                                                                                                                    |      |        |    |       |
| <b>TasMetaspaceEvictsRateHigh</b>                                                                                                                                                                                                           | 80   | 100    | 30 | FALSE |
| The rate at which 'evicts' are occurring is above the defined thresholds.<br><b>Index Type(s):</b> PerMetaspace                                                                                                                             |      |        |    |       |
| <b>TasMetaspaceExpireRateHigh</b>                                                                                                                                                                                                           | 80   | 100    | 30 | FALSE |
| The rate at which 'expires' are occurring is above the defined thresholds.<br><b>Index Type(s):</b> PerMetaspace                                                                                                                            |      |        |    |       |
| <b>TasMetaspaceGetRateHigh</b>                                                                                                                                                                                                              | 80   | 100    | 30 | FALSE |
| The rate at which 'gets' are occurring is above the defined thresholds.<br><b>Index Type(s):</b> PerMetaspace                                                                                                                               |      |        |    |       |
| <b>TasMetaspacePutRateHigh</b>                                                                                                                                                                                                              | 80   | 100    | 30 | FALSE |
| The rate at which 'puts' are occurring is above the defined thresholds.<br><b>Index Type(s):</b> PerMetaspace                                                                                                                               |      |        |    |       |
| <b>TasMetaspaceTakeRateHigh</b>                                                                                                                                                                                                             | 80   | 100    | 30 | FALSE |
| The rate at which 'takes' are occurring is above the defined thresholds.<br><b>Index Type(s):</b> PerMetaspace                                                                                                                              |      |        |    |       |
| <b>TasQueryDurationHigh</b>                                                                                                                                                                                                                 | 4    | 5      | 30 | FALSE |
| The query duration (in seconds) is above the defined threshold (in seconds).<br><b>Index Type(s):</b> PerSpace                                                                                                                              |      |        |    |       |
| <b>TasSpaceEntriesHigh</b>                                                                                                                                                                                                                  | 8000 | 100000 | 30 | FALSE |
| The number of objects inserted into the space is above the defined thresholds.<br><b>Index Type(s):</b> PerSpace                                                                                                                            |      |        |    |       |
| <b>TasSpaceEvictsRateHigh</b>                                                                                                                                                                                                               | 80   | 100    | 30 | FALSE |
| The rate at which 'evicts' are occurring is above the defined thresholds.<br><b>Index Type(s):</b> PerSpace                                                                                                                                 |      |        |    |       |

|                                                                                                                                               |     |     |    |       |
|-----------------------------------------------------------------------------------------------------------------------------------------------|-----|-----|----|-------|
| <b>TasSpaceExpireRateHigh</b><br>The rate at which 'expires' are occurring is above the defined thresholds.<br><b>Index Type(s):</b> PerSpace | 80  | 100 | 30 | FALSE |
| <b>TasSpaceGetRateHigh</b><br>The rate at which 'gets' are occurring is above the defined thresholds.<br><b>Index Type(s):</b> PerSpace       | 80  | 100 | 30 | FALSE |
| <b>TasSpacePutRateHigh</b><br>The rate at which 'puts' are occurring is above the defined thresholds.<br><b>Index Type(s):</b> PerSpace       | 80  | 100 | 30 | FALSE |
| <b>TasSpaceSeederCountLow</b><br>Not enough seeders are available.<br><b>Index Type(s):</b> PerSpace                                          | NaN | NaN | 30 | FALSE |
| <b>TasSpaceState</b><br>The state of the space is "not ready".<br><b>Index Type(s):</b> PerSpace                                              | NaN | NaN | 30 | FALSE |
| <b>TasSpaceTakeRateHigh</b><br>The rate at which 'takes' are occurring is above the defined thresholds.<br><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.

| Alert Name                                                                                                                                                                                                | WARN. LEVEL | ALARM LEVEL | DURATION | ENABLED |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|-------------|----------|---------|
| <b>TadAdapterDeltaErrorsHigh</b><br>The number of errors incurred by the adapter in last measurement interval is above the defined threshold.<br><b>Index Type(s):</b> PerAdapter                         | 1600        | 2000        | 30       | FALSE   |
| <b>TadAdapterExpired</b><br>The data from this adapter has not been updated since the last measurement interval. The data shown from this adapter is currently stale.<br><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 <b>rate at which concepts are written to the cache</b> 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/<br/>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 (<code>currentMsSentCount &gt; 0</code>), but none of the messages have been consumed within the defined warning and alert thresholds (<code>elapsedSinceLasAckInSec &gt; threshold</code>). Alert and warning thresholds are in seconds.</p> <p><b>Index Type(s):</b> PerConsumer: ID/<br/>PerServerConsumer: URL; ID</p> <p><b>Metric:</b> currentMsgSentCount,<br/>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/<br/>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/<br/>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/<br/>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><br/>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/<br/>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/<br/>PerServerQueue: URL; name</p> <p><b>Metric:</b> producerCount</p>                                           | 15 | 5  | 30 | TRUE  |

|                                                                                                                                                                                                                  |     |     |    |       |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----|-----|----|-------|
| <b>EmsServerAsyncDBSizeHigh</b><br>The size of the Async database, in bytes, for the EMS Server reached its maximum.<br><b>Index Type(s):</b> PerServer:URL<br><b>Metric:</b> asyncDBSize                        | 50  | 100 | 30 | FALSE |
| <b>EmsServerInboundDeltaHigh</b><br>The number of new incoming messages for the EMS Server has reached its maximum<br><b>Index Type(s):</b> PerServer<br><b>Metric:</b> DeltainboundMessageCount                 | 60  | 80  | 30 | FALSE |
| <b>EmsServerSyncDBSizeHigh</b><br>The size of the Sync database, in bytes, for the EMS Server reached its maximum.<br><b>Index Type(s):</b> PerServer:URL<br><b>Metric:</b> syncDBSize                           | 50  | 100 | 30 | FALSE |
| <b>EmsServerConnectionCountHigh</b><br>Alert is triggered when the number of connections to the server reaches the specified threshold.<br><b>Index Type(s):</b> PerServer:URL<br><b>Metric:</b> connectionCount | 60  | 80  | 30 | FALSE |
| <b>EmsServerInMsgRateHigh</b><br>The number of inbound messages on the server exceeded the specified threshold.<br><b>Index Type(s):</b> PerServer:URL<br><b>Metric:</b> inboundMessageRate                      | 2   | 80  | 30 | FALSE |
| <b>EmsServerMemUsedHigh</b><br>The percent memory used on the server exceeded the specified threshold.<br><b>Index Type(s):</b> PerServer:URL<br><b>Metric:</b> messageMemoryPct                                 | 60  | 80  | 30 | FALSE |
| <b>EmsServerNotStarted</b><br>The server state is empty. The server is not started.<br><b>Index Type(s):</b> PerServer:URL<br><b>Metric:</b> NotStarted                                                          | NaN | NaN | 30 | FALSE |
| <b>EmsServerOutMsgRateHigh</b><br>The number of outbound messages on the server exceeded the specified threshold.<br><b>Index Type(s):</b> PerServer:URL<br><b>Metric:</b> outboundMessageRate                   | 60  | 80  | 30 | FALSE |
| <b>EmsServerPendingMsgsHigh</b><br>The number of pending messages in the server queue exceeded the specified threshold.<br><b>Index Type(s):</b> PerServer:URL<br><b>Metric:</b> pendingMessageCount             | 60  | 80  | 30 | FALSE |

|                                                                                                                                                                                                                                                                                                                                         |     |     |    |       |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----|-----|----|-------|
| <b>EmsServerPendingMsgSizeHigh</b><br>The size, in KB, of the pending messages stored on this EMS Server reached its maximum.<br><b>Index Type(s):</b> PerServer:URL<br><b>Metric:</b> pendingMessageCount                                                                                                                              | 60  | 80  | 30 | FALSE |
| <b>EmsServerRouteState</b><br>One or more routes on the server are not active.<br><b>Index Type(s):</b> PerServer:URL<br><b>Metric:</b> Alert State                                                                                                                                                                                     | NaN | NaN | 30 | FALSE |
| <b>EmsServerStaleData</b><br>The server stopped receiving data.<br><b>Index Type(s):</b> PerServer:URL<br><b>Metric:</b> Expired                                                                                                                                                                                                        | NaN | NaN | 30 | FALSE |
| <b>EmsTopicConsumerIdleTimeHigh</b><br>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).<br><b>Index Type(s):</b> PerTopic; PerServerTopic<br><b>Metric:</b> ConsumerIdleTime | 60  | 80  | 30 | FALSE |
| <b>EmsTopicInboundDeltaHigh</b><br>The number of new incoming messages for the EMS Topic has reached its maximum.<br><b>Index Type(s):</b> PerTopic; PerServerTopic<br><b>Metric:</b> DeltainboundTotalMessages                                                                                                                         | 60  | 80  | 30 | FALSE |
| <b>EmsTopicMsgLatencyHigh</b><br>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.<br><b>Index Type(s):</b> PerServerTopic<br><b>Metric:</b> messageLatency      | 60  | 80  | 30 | FALSE |
| <b>EmsTopicProviderIdleTimeHigh</b><br>The topic idle time exceeded the specified threshold. A topic is idle when the number of inbound messages remains unchanged.<br><b>Index Type(s):</b> PerServerTopic: URL; name<br><b>Metric:</b> ProviderIdleTime                                                                               | 60  | 80  | 30 | FALSE |
| <b>EmsTopicsConsumerCountHigh</b><br>The number of consumers for the topic exceeded the specified threshold.<br><b>Index Type(s):</b> PerServerTopic: URL; name<br><b>Metric:</b> consumerCount                                                                                                                                         | 60  | 80  | 30 | FALSE |
| <b>EmsTopicsConsumerCountLow</b><br>The number of consumers for the topic is below the specified threshold.<br><b>Index Type(s):</b> PerServerTopic<br><b>Metric:</b> consumerCount                                                                                                                                                     | 60  | 80  | 30 | FALSE |

|                                                                                                                                                                                                                                               |    |    |    |       |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----|----|----|-------|
| <b>EmsTopicsInMsgRateHigh</b><br>The number of inbound messages for the topic exceeded the specified threshold.<br><b>Index Type(s):</b> PerServerTopic<br><b>Metric:</b> inboundMessageRate                                                  | 60 | 80 | 30 | FALSE |
| <b>EmsTopicsOutMsgRateHigh</b><br>The rate of outbound messages for the topic exceeded the specified threshold.<br><b>Index Type(s):</b> PerServerTopic<br><b>Metric:</b> outboundMessageRate                                                 | 60 | 80 | 30 | TRUE  |
| <b>EmsTopicsPendingMsgsHigh</b><br>The number of pending messages on the queue for the topic exceeded the specified threshold.<br><b>Index Type(s):</b> PerTopic<br><b>Metric:</b> pendingMessageCount                                        | 50 | 75 | 30 | FALSE |
| <b>EmsTopicsProducerCountHigh</b><br>The number of active producers for this topic exceeded the specified high threshold.<br><b>Index Type(s):</b> PerTopic/PerServerTopic<br><b>Metric:</b> producerCount                                    | 60 | 80 | 30 | TRUE  |
| <b>EmsTopicsProducerCountLow</b><br>The number of producers for the topic is below the specified threshold.<br><b>Index Type(s):</b> PerTopic/PerServerTopic<br><b>Metric:</b> producerCount                                                  | 60 | 80 | 30 | TRUE  |
| <b>EmsTopicsSubscriberCountHigh</b><br>The number of subscribers for the topic exceeded the specified threshold.<br><b>Index Type(s):</b> PerServerTopic<br><b>Metric:</b> subscriberCount                                                    | 50 | 75 | 30 | FALSE |
| <b>JvmCpuPercentHigh</b><br>The percent JVM CPU usage exceeded the specified threshold.<br><b>Index Type(s):</b> PerJVM<br><b>Metric:</b> CpuPercent                                                                                          | 30 | 40 | 30 | FALSE |
| <b>JvmGcDutyCycleHigh</b><br>The JVM Garbage Collection contains an item that exceeded the specified duty cycle threshold (the percent of time spent in Garbage Collection).<br><b>Index Type(s):</b> PerGC<br><b>Metric:</b> TimeUsedPercent | 50 | 75 | 30 | FALSE |
| <b>JvmMemoryUsedHigh</b><br>The percent JVM memory used exceeded the specified threshold.<br><b>Index Type(s):</b> PerJVM<br><b>Metric:</b> MemoryUsedPercent                                                                                 | 50 | 75 | 30 | FALSE |

|                                                                                                                  |     |     |    |       |
|------------------------------------------------------------------------------------------------------------------|-----|-----|----|-------|
| <b>JvmNotConnected</b><br>The JVM is not connected.<br><b>Index Type(s):</b> PerJVM<br><b>Metric:</b> Connected  | NaN | NaN | 30 | FALSE |
| <b>JvmStaleData</b><br>The JVM stopped receiving data.<br><b>Index Type(s):</b> PerJVM<br><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>TftIClientCpuTime</b><br>Executes a single warning and a single alarm if the CPU response time to the client exceeds the specified threshold.<br>Index Type: PerClient<br>Metric: Delta_CPU_TIME                                  | 160000        | 200000      | 30       | FALSE   |
| <b>TftIClientCpuUsage</b><br>Executes a single warning and a single alarm if the CPU usage by the client exceeds the specified threshold.<br>Index Type: PerClient<br>Metric:                                                        | 160000        | 200000      | 30       | FALSE   |
| <b>TftIClientExpired</b><br>Executes a single alert if the response time to the client exceeds the specified threshold.<br>Index Type: PerClient<br>Metric:                                                                          | NaN           | NaN         | 30       | FALSE   |
| <b>TftIClientMemory</b><br>Executes a single warning and a single alarm if the memory usage by the client exceeds the specified threshold.<br>Index Type: PerClient<br>Metric: PROCESS_RSS_KB                                        | 160000        | 200000      | 30       | FALSE   |
| <b>TftIClientMsgsRcvdRate</b><br>Executes a single warning and a single alarm if the number of messages received by the client per second exceeds the specified threshold.<br>Index Type: PerClient<br>Metric: RateMESSAGES_RECEIVED | 160000        | 200000      | 30       | FALSE   |

|                                                                                                                                                                                                                                                                     |        |        |    |       |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------|--------|----|-------|
| <b>TftIClientMsgsSentRate</b><br>Executes a single warning and a single alarm if the number of messages sent by the client per second exceeds the specified threshold.<br>Index Type: PerClient<br>Metric: RateMESSAGES_SENT                                        | 160000 | 200000 | 30 | FALSE |
| <b>TftIClientNotRunning</b><br>Executes a single if the client status is not "RUNNING".<br>Index Type: PerClient<br>Metric: Delayed Writes                                                                                                                          | NaN    | NaN    | 30 | FALSE |
| <b>TftIClientVirtualMemory</b><br>Executes a single warning and a single alarm if the virtual memory usage by the client exceeds the specified threshold.<br>Index Type: PerClient<br>Metric:                                                                       | 160000 | 200000 | 30 | FALSE |
| <b>TftIServerClientCount</b><br>Executes a single warning and a single alarm if the number of clients on the FTL server exceeds the specified threshold.<br>Index Type: PerServer<br>Metric:                                                                        | 160    | 200    | 30 | FALSE |
| <b>TftIServerCpuTime</b><br>Executes a single warning and a single alarm if the FTL server CPU response time exceeds the specified threshold.<br>Index Type: PerServer<br>Metric:                                                                                   | 160    | 200    | 30 | FALSE |
| <b>TftIServerCpuUsage</b><br>Executes a single warning and a single alarm if the FTL server CPU usage exceeds the specified threshold.<br>Index Type: PerServer<br>Metric:                                                                                          | 60     | 80     | 30 | FALSE |
| <b>TftIServerExpired</b><br>Executes a single warning and a single alarm if the FTL server response time exceeds the specified threshold.<br>Index Type: PerServer<br>Metric:                                                                                       | NaN    | NaN    | 30 | FALSE |
| <b>TftIServerInboxSendFaults</b><br>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.<br>Index Type: PerServer<br>Metric: SEND_TO_INBOX_FAILURES | 160    | 200    | 30 | FALSE |

|                                                                                                                                                                                                                                                         |     |     |    |       |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----|-----|----|-------|
| <b>TftIServerMemory</b><br>Executes a single warning and a single alarm if the FTL server memory usage exceeds the specified threshold.<br>Index Type: PerServer<br>Metric: PROCESS_RSS_KB                                                              | 160 | 200 | 30 | FALSE |
| <b>TftIServerOnBackup</b><br>Executes a single alert if the primary FTL server is down and now running on the backup FTL server.<br>Index Type: PerServer<br>Metric:                                                                                    | NaN | NaN | 30 | FALSE |
| <b>TftIServerSatelliteCount</b><br>Executes a single alert if the number of satellite servers is lower than expected.<br>Note: Set threshold to one less than number of deployed satellites.<br>Index Type: Response Time<br>Metric: Table_locks_waited | NaN | 5   | 30 | FALSE |
| <b>TftIServerVirtualMemory</b><br>Executes a single warning and a single alarm if the FTL server virtual memory usage exceeds the specified threshold.<br>Index Type: Response Time<br>Metric: Table_locks_waited                                       | 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.

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>UXRobotError</b>        | 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> .<br><br>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. |
| <b>UXRobotResponseSlow</b> | 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.                                                                                                                                                                                                                                                 |



|                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <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><br>The Host's CPU utilization is above the defined threshold.<br><b>Index Type(s):</b> PerVmHost        | 50          | 75          | 2        | TRUE    |
| <b>VmHostDiskBytesReadHigh</b><br>The disk read rate (kBytes/second) is above the defined thresholds.<br><b>Index Type(s):</b> PerVmHost | 1024        | 2048        | 2        | TRUE    |

|                                                                                                                                                                   |       |       |   |      |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------|-------|---|------|
| <b>VmHostDiskBytesWrittenHigh</b><br>The disk write rate (kBytes/second) is above the defined thresholds.<br><b>Index Type(s):</b> PerVmHost                      | 1024  | 2048  | 2 | TRUE |
| <b>VmwHostInBytesHigh</b><br>The inbound byte rate (KB/second) is above the defined thresholds.<br><b>Index Type(s):</b> PerVmHost                                | 1024  | 2048  | 2 | TRUE |
| <b>VmwHostInPktDropLossHigh</b><br>The percentage of inbound packets dropped is above the defined threshold.<br><b>Index Type(s):</b> PerVmHost                   | 1     | 3     | 2 | TRUE |
| <b>VmwHostInPktErrorLossHigh</b><br>The percentage of inbound packets discarded for any error is above the defined threshold.<br><b>Index Type(s):</b> PerVmHost  | 1     | 3     | 2 | TRUE |
| <b>VmwHostMemoryUsageHigh</b><br>The percentage memory utilization (used/configured) is above the defined threshold.<br><b>Index Type(s):</b> PerVmHost           | 70    | 80    | 2 | TRUE |
| <b>VmwHostOutBytesHigh</b><br>The outbound byte rate (KB/second) is above the defined threshold.<br><b>Index Type(s):</b> PerVmHost                               | 1024  | 2048  | 2 | TRUE |
| <b>VmwHostOutPktDropLossHigh</b><br>The percentage of outbound packets dropped is above the defined thresholds.<br><b>Index Type(s):</b> PerVmHost                | 1     | 3     | 2 | TRUE |
| <b>VmwHostOutPktErrorLossHigh</b><br>The percentage of inbound packets discarded for any error is above the defined threshold.<br><b>Index Type(s):</b> PerVmHost | 1     | 3     | 2 | TRUE |
| <b>VmwHostStatusBad</b><br>The overall status is not "green."<br><b>Index Type(s):</b> PerVmHost                                                                  | NaN   | NaN   | 2 | TRUE |
| <b>VmwHostSwapUsedHigh</b><br>The amount of swap space used by a host is above the defined thresholds.<br><b>Index Type(s):</b> PerVmHost                         | 10240 | 40960 | 2 | TRUE |
| <b>VmwVmCpuUtilizationHigh</b><br>The virtual machine CPU utilization is above the defined thresholds.<br><b>Index Type(s):</b> PerVm                             | 50    | 75    | 2 | TRUE |
| <b>VmwVmDiskBytesReadHigh</b><br>The disk read rate (KB/second) is above the defined thresholds.<br><b>Index Type(s):</b> PerVm                                   | 1024  | 2048  | 2 | TRUE |

|                                                                                                                                                                                                                                                                                       |      |      |    |      |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------|------|----|------|
| <b>VmwVmDiskBytesWrittenHigh</b><br>The disk write rate (KB/second) is above the defined thresholds.<br><b>Index Type(s):</b> PerVm                                                                                                                                                   | 1024 | 2048 | 2  | TRUE |
| <b>VmwVmDiskUsageHigh</b><br>The amount of disk space used by the virtual machine is above the defined threshold.<br><b>Index Type(s):</b> PerVm                                                                                                                                      | 85   | 95   | 30 | TRUE |
| <b>VmwVmInBytesHigh</b><br>The inbound byte rate (KB/second) is above the defined threshold.<br><b>Index Type(s):</b> PerVm                                                                                                                                                           | 1024 | 2048 | 2  | TRUE |
| <b>VmwVmInPktDropLossHigh</b><br>The percentage of inbound packet loss due to dropped packets is above the defined threshold.<br><b>Index Type(s):</b> PerVm                                                                                                                          | 1    | 3    | 30 | TRUE |
| <b>VmwVmMemoryUsageHigh</b><br>The percentage of memory utilization (active/ configured) is above the defined thresholds.<br><b>Index Type(s):</b> PerVm                                                                                                                              | 70   | 80   | 2  | TRUE |
| <b>VmwVmOutBytesHigh</b><br>The outbound byte rate is above the defined threshold.<br><b>Index Type(s):</b> PerVm                                                                                                                                                                     | 1024 | 2048 | 2  | TRUE |
| <b>VmwVmOutPktDropLossHigh</b><br>The percentage of outbound packet loss due to dropped packets on the virtual machine is above the defined threshold.<br><b>Index Type(s):</b> PerVm                                                                                                 | 1    | 3    | 2  | TRUE |
| <b>VmwVmStatusBad</b><br>The overall status for this virtual machine is not "green."<br><b>Index Type(s):</b> PerVm                                                                                                                                                                   | NaN  | NaN  | 2  | TRUE |
| <b>VmwVmSwapUsedHigh</b><br>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 the may occur in the guest operating system.<br><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 2.11.0 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 Jackson libraries are licensed under the Apache License Version 2.11.0

#### Apache License 2.0

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.

**O. 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 NONINFRINGEMENT. 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 NONINFRINGEMENT. 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!