TIBCO® RTView® for TIBCO BusinessEvents® User's Guide

Version 7.1

Copyright © 2014-2024 TIBCO Software Inc. All Rights Reserved.

All rights reserved. No part of this manual may be reproduced, in any form or by any means, without written permission from TIBCO Software Inc.

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

The information in this document is subject to change without notice and should not be construed as a commitment by the TIBCO Software Inc. TIBCO Software Inc. assumes no responsibility for any errors that may appear in this document.

The software described in this document is furnished under a license and may be used or copied only in accordance with the terms of such license.

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.

TIBCO Software Inc.
3303 Hillview Avenue
Palo Alto, CA 94304-1213

CUSTOMER SUPPORT
For an overview of TIBCO Support Services, and information about getting started with TIBCO Product Support, visit this site:

http://www.tibco.com/services/support/default.jsp

If you already have a valid maintenance or support contract, visit this site:

http://support.tibco.com

Entry to this site requires a username and password. If you do not have a username, you can request one.

TIBCO, TIBCO Hawk, TIBCO Rendezvous, and TIBCO Enterprise Message Service are trademarks and/or registered trademarks of TIBCO Software Inc. in the United States and other countries. They are mentioned in this document for identification purposes only.

TIBCO® RTView® for TIBCO BusinessEvents® contains components licensed under the Apache License Version 2.0.

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 © 1998-2023 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.

TIBCO® RTView® for TIBCO BusinessEvents® User's Guide

Preface **1**

Document Conventions 1

Third Party Notices 1

CHAPTER 1 Quick Start **2**

Prerequisites for Windows and UNIX/Linux Installations 2

UNIX/Linux Quick Start Steps 2

Windows Quick Start Steps 6

CHAPTER 2 Introduction to the Monitor **11**

Overview 11

System Requirements 11

Installation and Setup 11

Installation 11

Setup 12

Enabling Login in the Monitor 12

Limitations for TIBCO BusinessEvents 4.0 Installations 13

Upgrading the Monitor 13

Version 7.1 14

SNMP Notifications 14

Alert Notification 14

Admin Displays (HTML UI) 14

Project Directory Structure 15

Version 7.0 15

CHAPTER 3 Configuration **17**

Overview 17

Basic Steps 17

Configuring Data Collection 18

Configuring Historical Data 20

Defining the Storage of In Memory BEMON History 21

Defining Compaction Rules for BEMON 21

Defining Expiration and Deletion Duration for BEMON Metrics 22

Enabling/Disabling Storage of BEMON Historical Data 23

Defining a Prefix for All History Table Names for BEMON Metrics 24

Configuring Ports 25

Configuring the Database 27

Enable Monitoring of TIBCO APIX Engines ... 31
Enable Collection of Tomcat Historical Data ... 31
Configure Alert Notification ... 32
Alert Event Options ... 33
Alert Action Options ... 33
Run a Script ... 34
Execute Java Code ... 34
Customizing the Custom Command Handler ... 35
Add Email Notification ... 35
Send SNMP Trap ... 35
Run Command String ... 36
Conditional Filter ... 36
Configure High Availability ... 36
Overview of High Availability Architecture ... 37
Data Server High Availability ... 37
Display Server High Availability (Classic User Interface) ... 37
HTML User Interface High Availability ... 37
Historian High Availability ... 37
Requirements for Configuring High Availability ... 38
Steps for Configuring High Availability ... 38
Verifying the High Availability Configuration ... 39
Primary Data Server Log File ... 39
Backup Data Server Log File ... 39
Primary Historian Log File ... 40
Backup Historian Log File ... 40
Primary Display Server Log File ... 41
Backup Display Server Log File ... 41
CHAPTER 4 Deployment ... 42
Overview ... 42
Web Application Deployment ... 42
Windows ... 42
UNIX/Linux ... 43
RTView Server Components as Windows Services ... 44
Troubleshooting ... 47
Log Files ... 47

JAVA_HOME ... 47
Permissions ... 47
Network/DNS ... 47
Verify Data Received from Data Server ... 48
Verify Port Assignments ... 48
Sender/Receiver: Distributing the Load of Data Collection ... 48
Example ... 49
Setting Up the Sender/Receiver Configuration ... 50
Receiver Configuration ... 50
Collector Configuration ... 51
CHAPTER 5 Using the Monitor -- HTML Displays ... 53
Overview ... 53
TIBCO BusinessEvents Overview Display ... 53
Tables ... 54
Column Filtering ... 55
Multiple Column Sorting ... 57
Column Visibility ... 57
Column Reordering ... 58
Row Paging ... 58
Trend Graphs ... 58
Time Settings ... 58
Mouse-over ... 59
Log Scale ... 59
Title Bar Functionality ... 60
Export Report ... 60
Rate Calculation in TIBCO BusinessEvents Monitor ... 61
TIBCO BusinessEvents Monitor HTML Views/Displays ... 61
BE Clusters - HTML ... 62
TIBCO BE Clusters Table - HTML ... 62
TIBCO BE Clusters Heatmap - HTML ... 64
TIBCO BE Cluster Summary - HTML ... 65
BE Nodes - HTML ... 68
TIBCO BE Cluster Nodes Table - HTML ... 68
TIBCO BE Cluster Nodes Heatmap - HTML ... 70
TIBCO BE Inference Node Summary - HTML ... 72

TIBCO BE Storage Node Summary - HTML ... 76
BE Events - HTML ... 80
TIBCO BE Events Table - HTML ... 80
TIBCO BE Agent Event Summary - HTML ... 82
TIBCO BE Event Cache Hits Table - HTML ... 86
TIBCO BE Event Hit Summary - HTML ... 87
BE Concepts - HTML ... 90
TIBCO BE Concepts Table - HTML ... 90
TIBCO BE Concept Hit Summary - HTML ... 92
TIBCO BE Channel Status Table - HTML ... 95
TIBCO BE Inference Agents Table - HTML ... 97
TIBCO BE RTC Txn Manager Reports - HTML ... 99
Drilldowns ... 101
Alerts Table by Component ... 101
Alert Detail for Component ... 103
Alerts ... 104
Alerts Table ... 105
Admin ... 106
Alerts Administration ... 106
To set thresholds and enable an override alert: ... 107
Alert Overrides Administration ... 108
Alert Engine Admin ... 110
Disable Alert Engine ... 110
Enable Alert Engine ... 111
Alert Configuration for Component ... 111
Cache Table ... 112
CHAPTER 6 Using the Monitor ... 115
Overview ... 115
Monitor Main Display ... 115
Tables ... 116
Multiple Column Sorting ... 117
Column Visibility ... 117
Column Filtering ... 118
Column Locking ... 119
Column Reordering ... 120

Saving Settings ... 120
Revert Table Settings ... 120
Row Paging ... 120
Row Color Code ... 121
Row Keyboard Selection ... 121
Trend Graphs ... 122
Time Range ... 122
Mouse-over ... 122
Log Scale ... 123
Title Bar Functionality ... 123
Multiple Windows ... 124
Context Menu ... 124
Export Report ... 124
Rate Calculation in TIBCO BusinessEvents Monitor ... 125
TIBCO BusinessEvents Monitor Views/Displays ... 126
Clusters / Nodes View ... 126
Clusters ... 127
Cluster Summary ... 129
Cluster Nodes Table ... 132
Cluster Nodes Heatmap ... 134
Inference Node Summary ... 137
Storage Node Summary ... 142
Events / Concepts View ... 145
Agent Events ... 145
Agent Event Summary ... 147
Event Cache Hits ... 150
Event Hit Summary ... 152
Concept Cache Hits ... 155
Concept Hit Summary ... 156
Channels ... 159
All Inference Agents ... 161
All RTC Reports ... 163
Alert Views ... 165
Alert Detail Table ... 165
Administration ... 168

Alert Administration ... 168
Global Thresholds ... 168
Override Thresholds ... 169
Tabular Alert Administration ... 171
Setting Override Alerts ... 173
Alert Administration Audit ... 173
Metrics Administration ... 175
RTView Cache Tables ... 176
Agent Administration ... 178
Version Info ... 179
About ... 180
APPENDIX A Monitor Scripts ... 182
Scripts ... 182
rtvservers.dat ... 191
Single Configuration File ... 191
Multiple Configuration File ... 192
APPENDIX B Alert Definitions ... 193
APPENDIX C RTView Configuration Application ... 196
Accessing the RTView Configuration Application ... 196
Via URL ... 197
Via a Button in the Monitor ... 197
Projects Page ... 198
Server Configuration View ... 198
General ... 199
General Tab ... 199
Custom Properties Tab ... 200
Databases ... 201
Connections Tab ... 202
Alerts ... 203
Alerts Tab ... 203
History Tab ... 207
Security ... 208
SSL Credentials ... 208
Securing RTView JMX Ports ... 209
Secure with SSL ... 209

Secure with Username and Password ... 209
Secure Client and Receiver Ports with SSL ... 211
Data Server ... 211
Data Server Tab ... 211
Collector Tab ... 213
Historian ... 213
Display Server ... 214
Solution Package Configuration View ... 215
APPENDIX D Security Configuration ... 217
Introduction ... 217
Port Settings ... 218
Data Server ... 219
Display Server (thin client) ... 220
HTML UI ... 220
Data Collectors ... 221
Configuration Application ... 221
Configuration Files ... 221
Historian ... 221
Database ... 222
Application Servers ... 222
Monitored Components ... 223
TIBCO BusinessEvents ... 223
TIBCO BusinessWorks 5 ... 223
TIBCO BusinessWorks 6 ... 224
TIBCO EMS Server ... 224
TIBCO Hawk ... 224
RTView Manager ... 224
RTView Processes ... 225
Security Summary ... 226
Secure Installation Location - High Priority ... 226
Login and Servlet Authentication - High Priority ... 227
Application Server Security - High Priority ... 227
Secure Connections between RTView Processes - Medium-to-Low Priority* ... 227
Secure Connections to Monitored Components - Medium-to-Low Priority* ... 227
Secure Connections to Databases - Medium-to-Low Priority* ... 227

Preface

Welcome to the TIBCO® RTView® for TIBCO BusinessEvents® User's Guide

Document Conventions

This guide uses the following standard set of typographical conventions.

Convention	Meaning
italics	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:
	`amnesiac > enable`
	`amnesiac # configure terminal`

< >	Values that you specify appear in angle brackets: **interface <ipaddress>**

Third Party Notices

Please refer to the **LICENSES_thirdparty.txt** file from your product installation.

CHAPTER 1 Quick Start

This chapter is intended for those customers evaluating TIBCO® RTView® for TIBCO BusinessEvents® for purchase and describes the basic steps required to install, configure, and start the Monitor using default settings while using Eclipse Jetty (which is delivered with the Monitor) as the application server. The steps listed in this chapter represent only the basic flow needed to get the Monitor up and running. See Introduction to the Monitor, Configuration, and Deployment for additional configuration and deployment options/details.

This chapter contains:

- Prerequisites for Windows and UNIX/Linux Installations
- UNIX/Linux Quick Start Steps
- Windows Quick Start Steps

Prerequisites for Windows and UNIX/Linux Installations

- TIBCO BusinessEvents 4.0.x, 5.0.x, 5.1.x, 6.x
- If you are using TIBCO Businessworks Version 6, you must enable TIBCO Hawk or install the OSGI Plugin. See Enable Monitoring via TIBCO Hawk for Versions 5 and 6 and Enable Monitoring via OSGi Plugin for Version 6 for more information.
- If you are upgrading to a newer version and want to copy your existing configurations/properties into the new version, see Upgrading the Monitor for information.
- Java JDK 8, 9, 11, or 17
- Application Server (for example, Eclipse Jetty (delivered with the Monitor), or Tomcat 8.5+)

For a more complete list of system requirements, refer to the **README_sysreq.txt** file from your product installation.

UNIX/Linux Quick Start Steps

1. Download **TIB_rtview-be_<version>.zip** to your local UNIX/Linux server.

   Note: If using UNIX, do not include spaces in your installation directory path. The **start_server.sh** and **stop_server.sh** scripts will not function properly if spaces are included in the installation directory path.

2. Extract the files:

   unzip -a TIB_rtview-be_<version>.zip

   If correctly installed, you should see the TIB_rtview-be directory with an rtvapm sub-directory, which should include tbemon as a subdirectory.

3. Set the **JAVA_HOME** environment variable to point to your Java installation. For example:

   export JAVA_HOME=/opt/Java/jdk1.9.0

4. Navigate to the **TIB_rtview-be** directory and type:

start_server.sh

5. Open a browser and type the following URL to open the RTView Configuration Application:
   http://localhost:3270/rtview-tbemon-rtvadmin
   Use rtvadmin/rtvadmin for the username/password.
   The RTView Configuration Application displays.

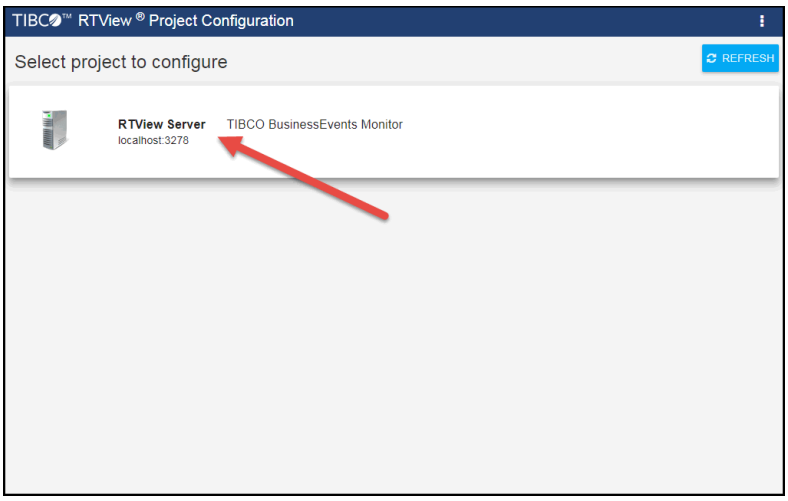

6. Select the **RTView Server** (**TIBCO BusinessEvents Monitor**) project, and then select **TIBCO BusinessEvents** under **Solution Package Configuration**.

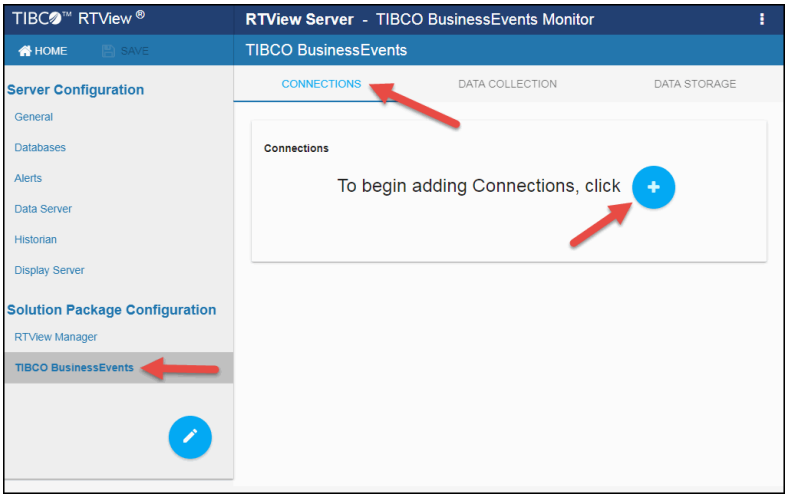

7. On the **CONNECTIONS** tab, click the + icon to add connections to your BEMON engines.
   The **Add Connection** dialog displays.

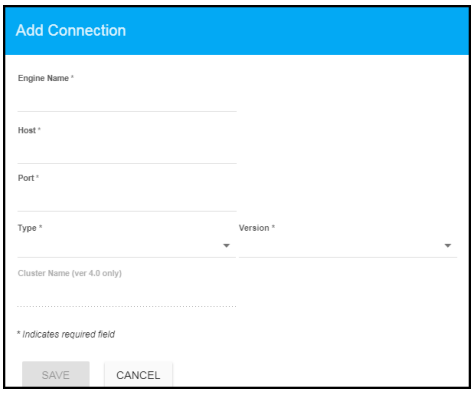

8. Enter the **Engine Name**, the **Host**, the **Type**, the **Version**, and the **Cluster Name** (if Version selected was 4.0) for the engine to which you want to connect (to enable the Monitor to collect data) where:

   Engine Name: the data source connection name used by the Monitor (it is not related to any TIBCO BusinessEvents configuration). Choose a descriptive name as the name appears in the Monitor displays for end-users. It should match the name specified in the first property.

   Host: resolves to the address of the system where the TIBCO BusinessEvents engine is running. Specify a unique hostname/port combination for each JMX connection. The hostname can be an IP address or a name that is resolvable via DNS or other network name resolution method used on the host.

   Port: the TCP port number assigned to the agent for monitoring via JMX. This port number is usually set in the engine's **.tra** file (**java.property.be.engine.jmx.connector.port=%jmx_port%**) and on the command line (**--propVar jmx_port=58700**).

   Type: Select the BusinessEvents engine type (**Cache** or **Inference**).

   Version: Select the BusinessEvents Version for the engine.

   Cluster Name: Specify the name of the cluster in which the engine resides. This field is only active if you selected version **4.0** in the **Version** field, and is intended to provide a solution if you have BusinessEvents nodes running on Java version 1.6.0_30 or earlier.

   NOTE: The BusinessEvents Solution Package requires Java version 1.6.0_31 or later, as it depends on JMX support for wild-carded queries to auto-discover BE cluster names.

9. click **Save**.

   Your newly created engine connection displays in the Connections tab.

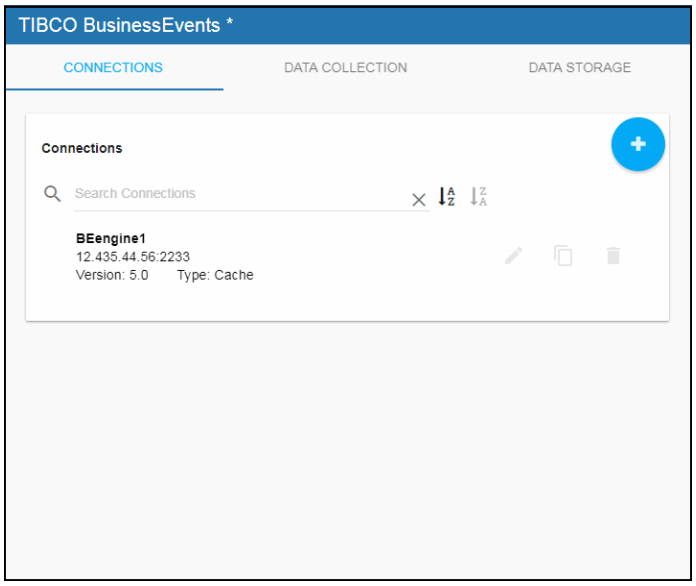

10. Optionally, if you are using the Classic (non-HTML) displays and want to use the white style sheet instead of the dark style sheet, and also want to enable Cross Platform Fonts, navigate to the **Display Server** option under **Server Configuration** and enable the **Use Light Stylesheet** and **Enable Cross Platform Fonts** toggles.

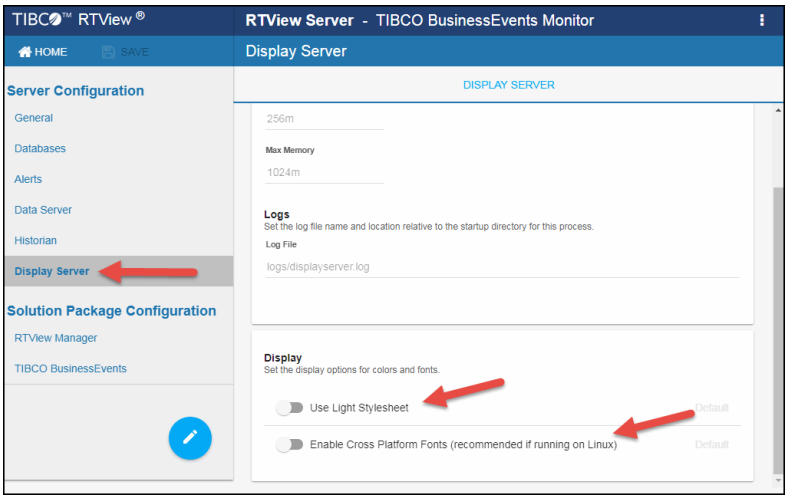

11. Click the **SAVE** button in the RTView Configuration Application to save your changes.

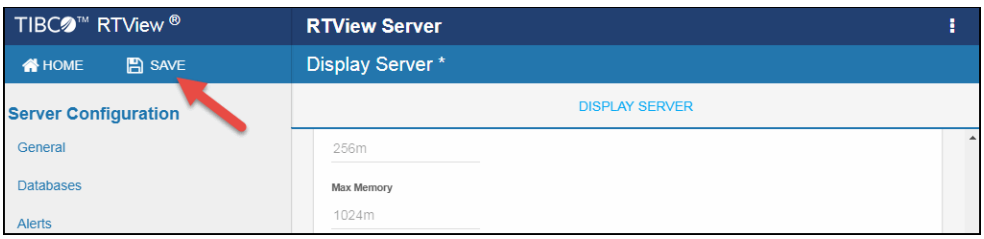

12. Click the RESTART SERVERS button (which displays in the upper right-hand corner after clicking **SAVE**) to apply the changes made in the RTView Configuration Application.

    The Project Page displays with a spinning icon. Once the data server has restarted, the spinning icon disappears and you can click your project and resume making changes (if desired).

13. You can check the log files in the **projects/rtview-server/logs** directory for errors.
14. You can verify that your caches are collecting data by accessing the following URL:

    http(s)://localhost:3270/common

    The RTView Cache Viewer application displays, which allows you to view the details for the caches that are collecting data.

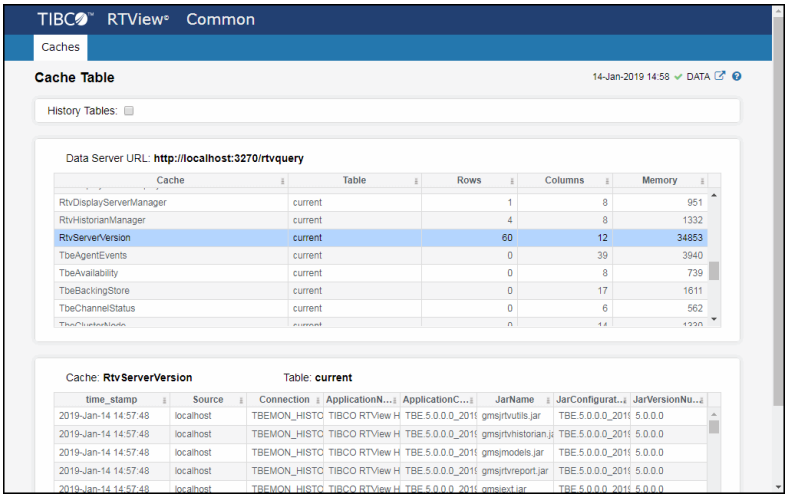

15. To view the monitor with HTML displays, open a browser and navigate to **http://-localhost:3270/rtview-tbemon** while using **rtvadmin**/**rtvadmin** as the login/-password.

    To view the monitor with Classic displays, open a browser and navigate to **http://localhost:3270/rtview-tbemon-classic** while using **rtvadmin**/**rtvadmin** as the login/password (if Login is enabled).

Windows Quick Start Steps

1. Download **TIB_rtview-be_<version>.zip** to your local Windows server.
2. Extract the files in **TIB_rtview-be_<version>.zip** using right mouse-click >"**Extract All...**"

   If correctly installed, you should see the TIB_rtview-be directory with an rtvapm sub-directory, which should include tbemon as a subdirectory.

3. Set the **JAVA_HOME** environment variable to point to your Java installation. For example:

   set JAVA_HOME=C:\Java\jdk1.9.0

4. Navigate to **TIB_rtview-be** directory and type:

   start_server.bat (in a command window)

   or

   Double-click **start_server.bat** in Windows

5. Open a browser and type the following URL to open the RTView Configuration Application:
   http://localhost:3270/rtview-tbemon-rtvadmin
   Use rtvadmin/rtvadmin for the username/password.
   The RTView Configuration Application displays.

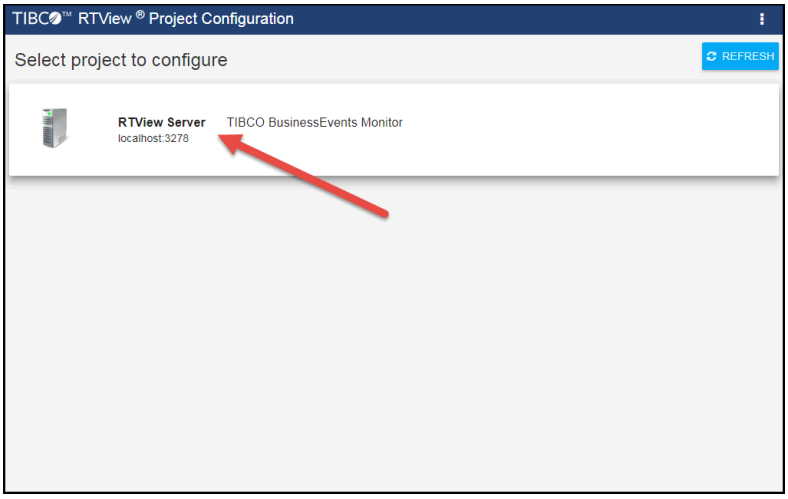

6. Select the **RTView Server** (**TIBCO BusinessEvents Monitor**) project, and then select **TIBCO BusinessEvents** under **Solution Package Configuration**.

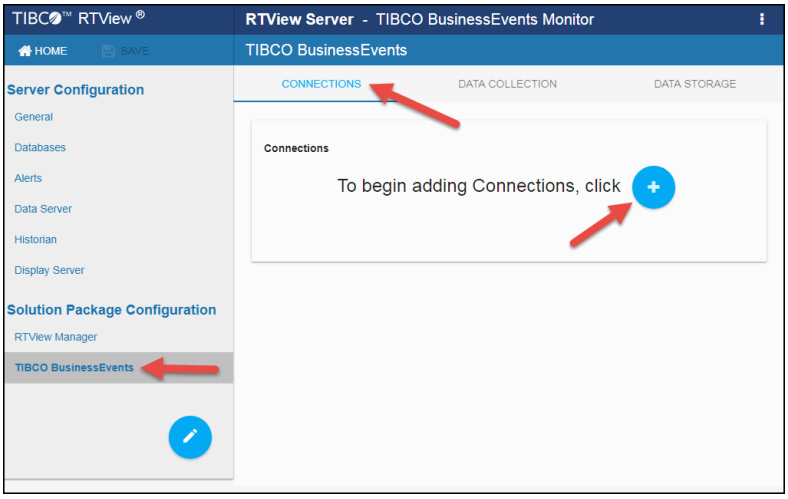

7. On the **CONNECTIONS** tab, click the + icon to add connections to your BEMON engines.
   The **Add Connection** dialog displays.

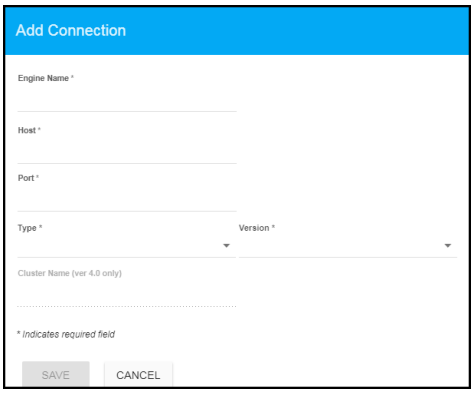

8. Enter the **Engine Name**, the **Host**, the **Type**, the **Version**, and the **Cluster Name** (if Version selected was 4.0) for the engine to which you want to connect (to enable the Monitor to collect data) where:

   Engine Name: the data source connection name used by the Monitor (it is not related to any TIBCO BusinessEvents configuration). Choose a descriptive name as the name appears in the Monitor displays for end-users. It should match the name specified in the first property.

   Host: resolves to the address of the system where the TIBCO BusinessEvents engine is running. Specify a unique hostname/port combination for each JMX connection. The hostname can be an IP address or a name that is resolvable via DNS or other network name resolution method used on the host.

   Port: the TCP port number assigned to the agent for monitoring via JMX. This port number is usually set in the engine's **.tra** file (**java.property.be.engine.jmx.connector.port=%jmx_port%**) and on the command line (**--propVar jmx_port=58700**).

   Type: Select the BusinessEvents engine type (**Cache** or **Inference**).

   Version: Select the BusinessEvents Version for the engine.

   Cluster Name: Specify the name of the cluster in which the engine resides. This field is only active if you selected version **4.0** in the **Version** field, and is intended to provide a solution if you have BusinessEvents nodes running on Java version 1.6.0_30 or earlier.

   NOTE: The BusinessEvents Solution Package requires Java version 1.6.0_31 or later, as it depends on JMX support for wild-carded queries to auto-discover BE cluster names.

9. click **Save**.

   Your newly created engine connection displays in the Connections tab.

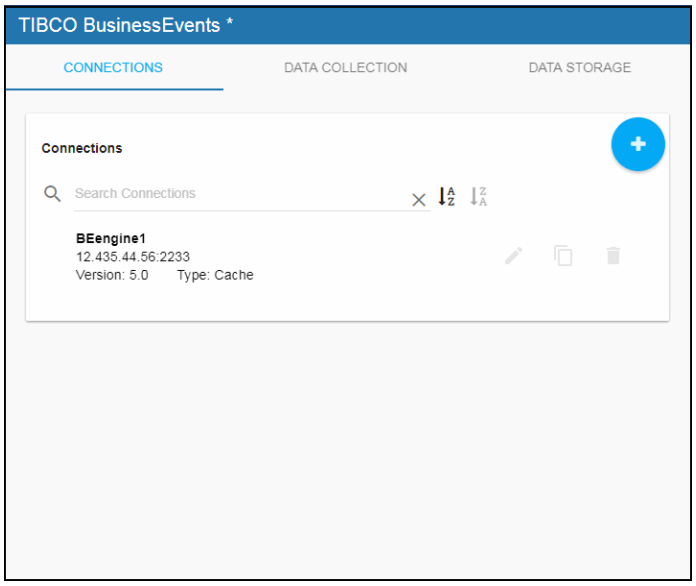

10. Optionally, if you are using the Classic (non-HTML) displays and want to use the white style sheet instead of the dark style sheet, and also want to enable Cross Platform Fonts, navigate to the **Display Server** option under **Server Configuration** and enable the **Use Light Stylesheet** and **Enable Cross Platform Fonts** toggles.

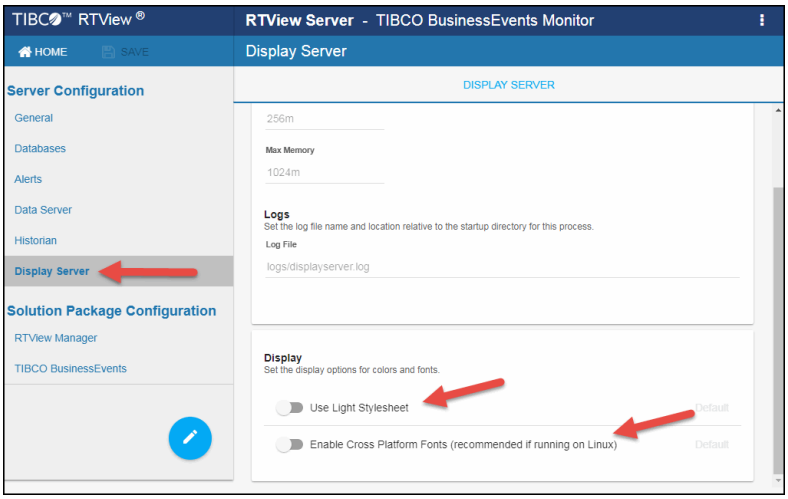

11. Click the **SAVE** button in the RTView Configuration Application to save your changes.

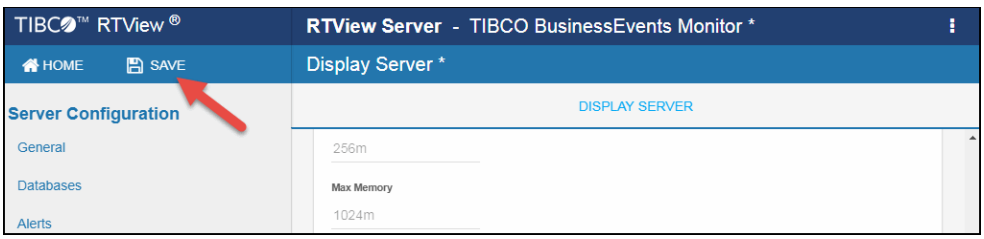

12. Click the RESTART SERVERS button (which displays in the upper right-hand corner after clicking **SAVE**) to apply the changes made in the RTView Configuration Application.

    The Project Page displays with a spinning icon. Once the data server has restarted, the message disappears and you can click your project and resume making changes (if desired).

13. You can check the log files in the **projects\rtview-server\logs** directory for errors.
14. You can verify that your caches are collecting data by accessing the following URL:

    http(s)://localhost:3270/common

    The RTView Cache Viewer application displays, which allows you to view the details for the caches that are collecting data.

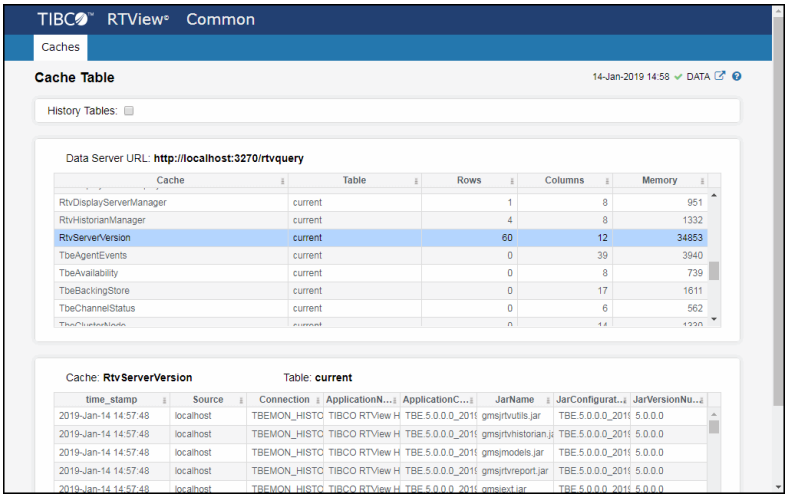

15. To view the monitor with HTML displays, open a browser and navigate to **http://-localhost:3270/rtview-tbemon** while using **rtvadmin**/**rtvadmin** as the login/-password.

    To view the monitor with Classic displays, open a browser and navigate to **http://localhost:3270/rtview-tbemon-classic** while using **rtvadmin**/**rtvadmin** as the login/password (if Login is enabled).

See Configuration, Deployment, and Using the Monitor for additional information about the configuration options, deployment options, and the various displays available in the Monitor.

CHAPTER 2 Introduction to the Monitor

This section contains the following:

- "Overview" on page 11
- "System Requirements" on page 11
- "Installation and Setup" on page 11
- "Upgrading the Monitor" on page 13

Overview

The Monitor provides information about how TIBCO BusinessEvents clusters are configured and performing, presents historical data detailing rule execution times per inference node, heap and table sizes for storage nodes, and event, concept and channels statistics. Pre-configured alert conditions provide early warning when any of these gathered performance metrics indicate a situation which is nearing a critical state.

The Monitor can help to diagnose several critical conditions relevant to the health of TIBCO BusinessEvents, including:

- events flooding into the system at much higher-than-expected rates.
- rules firing at a much higher rate than expected causing CPU usage to spike.
- the backing store running inefficiently.
- BusinessEvents concepts being created at a much higher rate than expected causing evaluation or re-evaluation of rules.

System Requirements

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.

Installation and Setup

This section describes how to install the BusinessEvents Monitor. This section includes the following:

- Installation
- Setup
- Limitations for TIBCO BusinessEvents 4.0 Installations

Installation

TIBCO BusinessEvents Monitor can be used as a standalone monitoring system for technical support teams. To install TIBCO BusinessEvents Monitor, download the **TIB_rtview-be_**

<version>.zip archive, and unzip the **TIB_rtview-be_<version>.zip** file into a directory of your choosing. See "Quick Start" on page 2 for more information.

File Extraction Considerations

On Windows systems, using the extraction wizard of some compression utilities might result in an extra top-level directory level based on the name of the **.zip** file. The additional directory is not needed because the **.zip** files already contain the **rtvapm** top-level directory. This extra directory must be removed before clicking the **Next** button that performs the final decompression.

To convert text files on UNIX/Linux systems to the native format, use the **-a** option with unzip to properly extract text files. Then, to fix execution permissions for all ***.sh** scripts, go to the **rtvapm** directory and execute:

. ./rtvapm_init.sh

Setup

This section describes how to setup your system for the Monitor.

- The memory setting for BusinessEvents projects must be set to **Cache**. The "in memory" memory management setting does not expose the MBeans queried by the Monitor. This means that your project must use a cluster of both inference and cache agents. Inference-only configurations are not supported, since the JMX MBean data is not available.
- Your BusinessEvents (inference and cache) engines must be JMX-enabled. For example, to start a simple BusinessEvents cluster on the command line, type:

start be-engine --propFile be-engine.tra -n TestCache --propVar jmx_port=58700 -c yourBeProjectCDD.cdd -u cache yourBeProjectEAR.ear

be-engine --propFile be-engine.tra -n TestInf --propVar jmx_port=58701 -c yourBeProjectCDD.cdd -u default yourBeProjectEAR.ear

Enabling Login in the Monitor

> **Note:** The following only enables login in the Display Server (Classic) version of the User Interface. The following steps do not enable login for the HTML User Interface, which has login enabled by default.

By default, "Login" is disabled forTIBCO® RTView® for TIBCO BusinessEvents®, which means that logging in will not be required when initially accessing TIBCO RTView for TIBCO BusinessEvents. To enable RTView Role Based Security, follow the instructions below. This will enable the following users and their associated passwords:

rtvadmin (password: rtvadmin)

admin (password: admin)

rtvuser (password: rtvuser)

user (password: user)

To define your own users and roles or to integrate with LDAP or other user and security management systems, see Role-based Security in the RTView Core User's Guide at https://slcorp.atlassian.net/wiki/spaces/RCS/pages/781516847/RTView+Documentation.

Note: This only enables login in the Display Server version of the User Interface. It does not enable login for the HTML User Interface.

To enable Login, perform the following steps:

1. Navigate to **TIB_rtview-bw/projects/rtview-server**.
2. Extract the **rtvdisplay.properties** file from the **rtview-tbemon-classic.war** file by typing:

   jar -xf rtview-tbemon-classic.war WEB-INF/classes/gmsjsp/rtvdisplay.properties

   The **WEB-INF/classes/gmsjsp** directory path is created, which contains the **rtvdisplay.properties** file.
3. Edit the **WEB-INF/classes/gmsjsp/rtvdisplay.properties** file, set **LoginEnabled=true**, and save your changes.
4. Update the **rtview-tbemon-classic.war** file with your changes by typing:

   jar -uf rtview-tbemon-classic.war WEB-INF\classes\gmsjsp\rtvdisplay.properties
5. If you are using Tomcat as your application server, copy the **TIB_rtview-bw/projects/rtview-server/rtview-tbemon-classic.war** file to the Tomcat **webapps** directory. If you are using Eclipse Jetty as your application server (which is delivered with TIBCO RTView for TIBCO BusinessEvents), there are no further steps.

Limitations for TIBCO BusinessEvents 4.0 Installations

Applications using JMX to monitor applications on a server with the server's firewall enabled might experience connection problems. The JMX protocol allows initial contact on a known port, but subsequent communications might occur over a second randomly chosen port. Version 5 of TIBCO BusinessEvents has a fix that allows the follow-on communications to occur over the same port. However, BusinessEvents version 4.0 does not have this fix. BusinessEvents 4.0 installations should use a local agent to push the necessary MBean data to the central RTView Data Server, or use a "premain agent" as described here:

https://blogs.oracle.com/jmxetc/entry/connecting_through_firewall_using_jmx

Upgrading the Monitor

This section describes the steps necessary to upgrade existing Monitor applications.

Follow the steps for each version between the version you are upgrading from and the version to which you are upgrading.

- Version 7.1
- Version 7.0

Version 7.1

Upgrading projects to Version 7.1 from the previous version requires the following:

Log4j2

The syntax used in a Log4j properties file was changed completely by Apache in version 2. The **sl.log4j.properties** file distributed with RTView has been changed to use the version 2 syntax. If, in previous RTView versions, you customized that file or specified your own custom log4j properties file (e.g. using the "log4jprops" option), you'll need to remake those customization using the version 2 syntax.

Customization changes should be made to the copy of **sl.log4j.properties** in the **projects** directory, instead of the copy under **rtvapm\common\conf** in order to make it easier toupgrade to future releases.

Note that the default logging behavior has been changed: In this release by default messages are appended to the existing **logs/X.log** file (where X is "dataserver", or "historian", etc depending on the name of the server) until it reaches a size of 50MB. Then it is renamed to X.log.N (where N = 1 - 9) and a new empty X.log file is created. So at any time the logs directory may contain X.log (newest, up to 50MB in size), and X.log.N where N = 1 - 9, each approx 50 MB, where 1 is the oldest and 9 is the newest. Once N = 9 is reached, on the next rollover X.log.1 is deleted and each remaining X.log.N is renamed to X.log.N-1.

SNMP Notifications

If you are upgrading from a previous release that sent SNMP notifications, you need to update the MIB in your SNMP receiver. The MIB definition in **rtvapm\common\lib\SL-RTVIEW-EM-MIB.txt** has changed to include a new field for this.

Alert Notification

In previous releases, any notification other than executing the **my_alert_actions** scripts had to be configured on the **CUSTOM PROPERTIES** tab of the RTView Configuration Application. If you are running a project configured with a previous release that has alert notifications defined in the **CUSTOM PROPERTIES** tab, they will continue to work as they did before with no changes. However, if you want to use the new **Alert Notification** configuration page, you will need to first delete the **CUSTOM PROPERTIES** for alert notifications, then recreate them in the **Alert Notification** page. Otherwise, the **CUSTOM PROPERTIES** will override the settings in the **Alert Notification** page.

Admin Displays (HTML UI)

For improved security the following displays have been moved from **assets/packages/common** to **assets/packages/admin**:

- Alert Administration (rtv_alerts_admin_table.html)
- Alert Overrides Admin (rtv_alerts_admin_overrides.html)
- Component Alert Configuration (rtv_alerts_admin_detail.html)

Any existing browser bookmarks to those displays should be updated or recreated.

Project Directory Structure

1. Copy ***.properties** and ***.properties.json** from your old project directory to the **projects\rtview-server** directory.

2. If you modified the **rtvservers.dat** in your old project, make the same changes to **projects\rtview-server\rtvservers.dat**.

3. If you are not using jetty, deploy the **.war** files from **projects\rtview-server** to your application server.

4. If you are using Jetty, open **project.properties** and **project.properties.json** in a text editor and replace **../../webapps/** with **./** in all war file paths in both files.

5. Use the top level **start**/**stop**/**status** scripts to start, stop and check the status of your monitor processes.

Version 7.0

Property File Handling and the RTView Configuration Application

Property file handling has been modified in order to support the "RTView Configuration Application". Existing applications will continue to work as before with no changes. However, customers should be aware of the following if they want to merge their old properties into the new version.

In previous releases, each sample project defined its own ports, sender target, and server identification properties. These properties have been removed from the sample properties and should be defined in the project properties instead. Sample projects have been updated to include these properties. Upward compatibility support is included for projects created prior to 6.8. In 6.8, the **rtview.properties** files in all sample projects were replaced with **project.properties** files. Any project with an **rtview.properties** file is recognized as a project created with a previous release. In that case, RTView will automatically read in the old ports, sender target, and server identification properties for all versions in the **rtview.properties** file. Therefore, projects created with previous versions will continue to run with no modifications. However, projects containing an **rtview.properties** file cannot be configured using the new RTView Configuration Application.

There are three options for merging properties from a previous version into the new version:

- The first option is to use the RTView Configuration Application to reapply the settings. If you have a lot of connections, this is not very practical but, if you only have a few, it could be worthwhile since you will be able to use the RTView Configuration Application for everything moving forward. To do this, run the new version, open the RTView Configuration Application, and apply all configurations that were part of your previous project. See "Configuration" for more information on how to configure your project using the RTView Configuration Application.
- The second option is to use your old properties files instead of the RTView Configuration Application. This has the downside that you cannot use the RTView Configuration Application moving forward. To do this, do the following in the sample project directory

1. Copy the properties files from the old project directory into the new project directory, including the old **rtview.properties**.

2. Remove the **project.properties** and **project.properties.json** from the new project directory.

3. Add your properties files to the appropriate lines in **servers\rtvservers.dat**.

   - The third option for applying solution package server settings from a previous version is a combination of the above, which has the benefit of allowing you to use the RTView Configuration Application without having to re-enter all of your connections.

1. Copy the properties files from the old project directory into the new project directory. Do NOT copy the old **rtview.properties** into the new project directory.

2. Edit the properties file you just copied over to comment out or remove all non-connection properties.

3. Run the new project and use the RTView Configuration Application to apply all settings from your previous project except connections. See "Configuration" for more information on how to configure your solution package servers using the RTView Configuration Application.

4. Add your properties files from step 2 to the appropriate lines in **servers\rtvservers.dat**.

5. Moving forward, new connections can be added via the RTView Configuration Application or by hand editing the properties file from step 2, whichever is more convenient. However, only connections added via the RTView Configuration Application will be editable in the RTView Configuration Application.

Note: The following files are read and written by the Configuration Application and should never be manually edited: **project.properties** and **project.properties.json**.

CHAPTER 3 Configuration

This section describes how to configure the Monitor components. This section includes:

- "Overview" on page 17
- "Configuring Data Collection" on page 18
- "Configuring Ports" on page 25
- "Configuring the Database" on page 27
- "Enable Monitoring of TIBCO APIX Engines" on page 31
- "Enable Collection of Tomcat Historical Data" on page 31
- "Configuring Alert Notification" on page 1
- "Configure High Availability " on page 36

Most of the configuration is completed in the RTView Configuration Application. See RTView Configuration Application and Quick Start for more information.

Overview

This section describes how to configure the Monitor as a standalone application.

Basic Steps

Some of the configuration steps described here are required (where noted) and others are optional.

- Step 1 (required): "Configuring Data Collection" on page 18. Define the TIBCO Servers and destinations to be monitored and create your project directory. This step must be performed before running any deployment of the Monitor.
- Step 2 (optional): "Configuring Ports" on page 25. Configure ports for the RTView Servers.
- Step 3 (optional): "Configuring the Database" on page 27. Configure a production database.
- Step 4 (optional): "Enable Monitoring of TIBCO APIX Engines" on page 31. Set up monitoring of TIBCO APIX engines in TIBCO.
- Step 5 (optional): "Enable Collection of Tomcat Historical Data" on page 31. Enable collection of Tomcat Historical data in the RTView Configuration Application.
- Step 6 (optional): "Configuring Alert Notification" on page 1. Configure alerts to execute an automated action (for example, to send an email alert).
- Step 7 (optional): "Configure High Availability " on page 36. Define redundant system components to mitigate a single point of failure within TIBCO BusinessEvents Monitor.

Assumptions

This document assumes that:

- you installed the Monitor per instructions in "Installation and Setup" on page 11.

Configuring Data Collection

This section describes how to configure data collection for TIBCO BusinessEvents Monitor. You configure the Data Servers by defining data source connections for each TIBCO BusinessEvents engine that you want to monitor in the RTView Configuration Application. There are two agent types that you can configure: Inference Agents and Cache Agents.

> **Note:** Your BusinessEvents project must include a cache agent as well as inference agents. TIBCO does not expose the management MBeans for inference agent only configurations.

Before You Begin:

- Verify that your BusinessEvents engines are JMX-enabled.
- For each engine to be monitored, obtain the JMX port assigned to that engine.

To Configure the Data Connection

Use the RTView Configuration Application to configure your data collection:

1. Navigate to **Solution Package Configuration** > **TIBCO BusinessEvents** > **CONNECTIONS** tab.

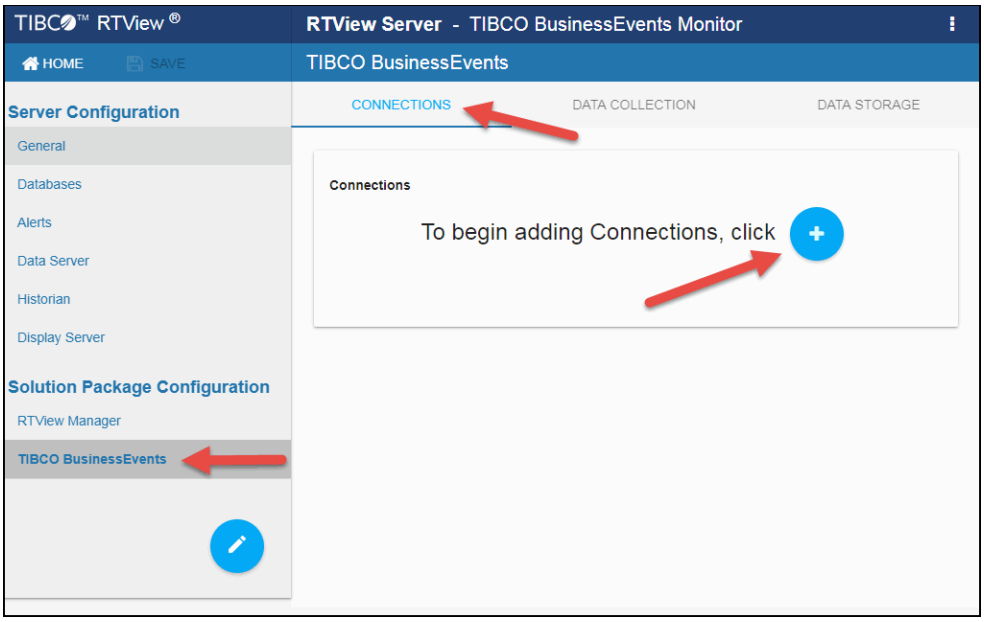

2. On the **CONNECTIONS** tab, click the icon to add connections to your BEMON engines.

   The **Add Connection** dialog displays.

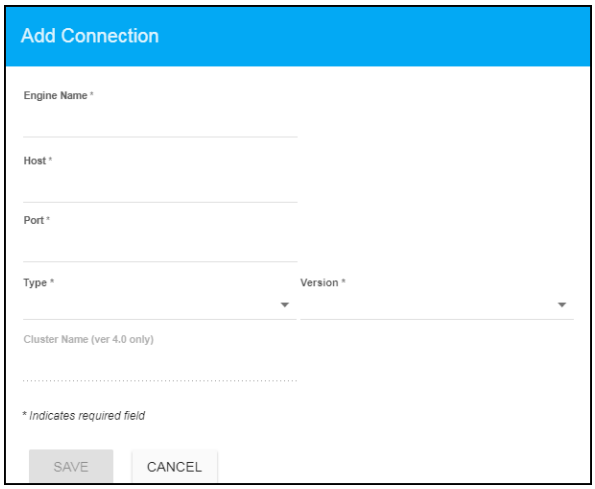

3. Enter the **Engine Name**, the **Host**, the **Type**, the **Version**, and the **Cluster Name** (if Version selected was 4.0) for the engine to which you want to connect (to enable the Monitor to collect data) where:

Engine Name: the data source connection name used by the Monitor (it is not related to any TIBCO BusinessEvents configuration). Choose a descriptive name as the name appears in the Monitor displays for end-users. It should match the name specified in the first property.

Host: resolves to the address of the system where the TIBCO BusinessEvents engine is running. Specify a unique hostname/port combination for each JMX connection. The hostname can be an IP address or a name that is resolvable via DNS or other network name resolution method used on the host.

Port: the TCP port number assigned to the agent for monitoring via JMX. This port number is usually set in the engine's **.tra** file (**java.property.be.engine.jmx.connector.port=%jmx_port%**) and on the command line (**--propVar jmx_port=58700**).

Type: Select the BusinessEvents engine type (**Cache** or **Inference**).

Version: Select the BusinessEvents Version for the engine.

Cluster Name: Specify the name of the cluster in which the engine resides. This field is only active if you selected version **4.0** in the **Version** field, and is intended to provide a solution if you have BusinessEvents nodes running on Java version 1.6.0_30 or earlier.

NOTE: The BusinessEvents Solution Package requires Java version 1.6.0_31 or later, as it depends on JMX support for wild-carded queries to auto-discover BE cluster names.

4. click **Save**.

Your newly created engine connection displays in the Connections tab.

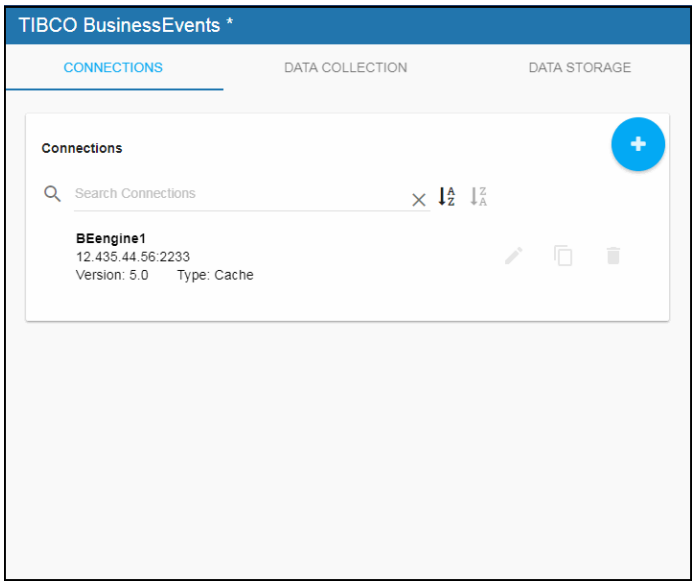

5. Repeat steps 2-4 to create any additional connections to BusinessEvents engines.

6. If you want to modify the default values for the update rates for the BusinessEvents caches, you can update the default polling rates in RTView Configuration Application > **DATA COLLECTION**> **Poll Rates**.

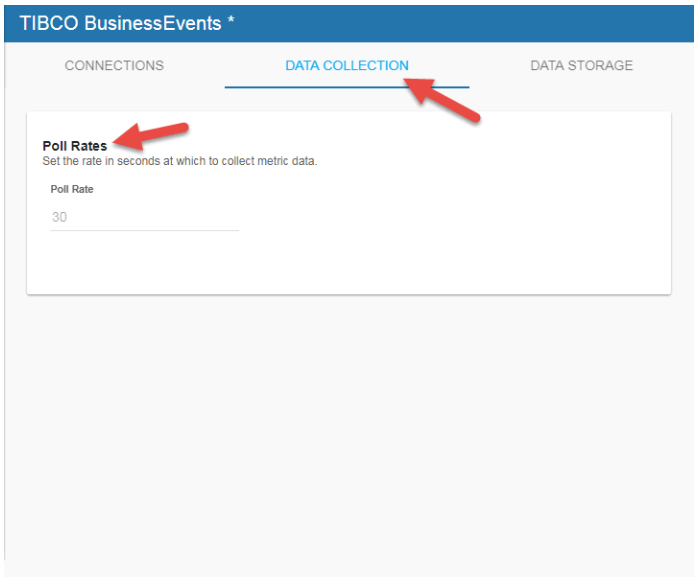

Configuring Historical Data

You can specify the number of history rows to store in memory, the compaction rules, the duration before metrics are expired and deleted, and the different types of metrics that you want the Historian to store in the **Data Storage** tab in the RTView Configuration Application. This section contains the following:

- Defining the Storage of In Memory BEMON History

- Defining Compaction Rules for BEMON
- Defining Expiration and Deletion Duration for BEMON Metrics
- Enabling/Disabling Storage of BEMON Historical Data
- Defining a Prefix for All History Table Names for BEMON Metrics

Defining the Storage of In Memory BEMON History

You can modify the maximum number of history rows to store in memory in the Data Storage tab. The **History Rows** property defines the maximum number of rows to store for the TbeInferenceAgent, TbeObjectTable, TbeBackingStore, TbeNodeEvents, TbeAgentEvents, TbeNodeConcepts, and TbeClusterSummary caches. The default settings for **History Rows** is 50,000. To update the default settings:

1. Navigate to **Solution Package Configuration** > **TIBCO BusinessEvents** > **DATA STORAGE** tab.
2. In the **Size** region, click the **History Rows** field and specify the desired number of rows.

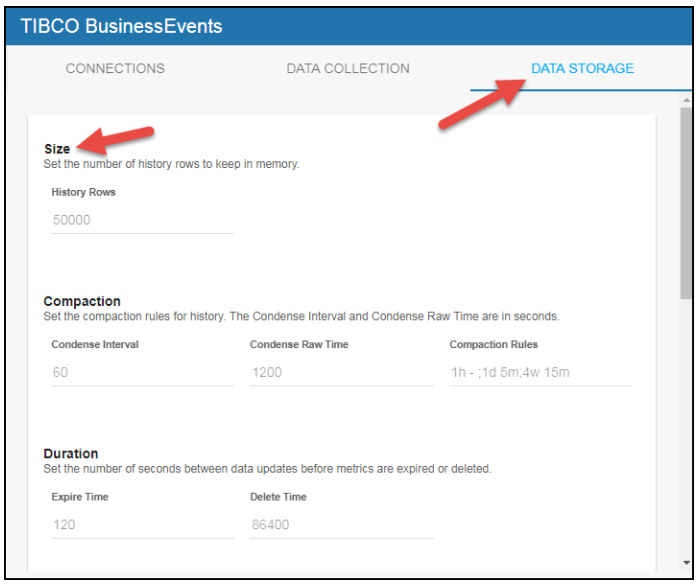

Defining Compaction Rules for BEMON

Data compaction, essentially, is taking large quantities of data and condensing it using a defined rule so that you store a reasonably sized sample of data instead of all of your data, thus preventing you from potentially overloading your database. The available fields are:

- **Condense Interval** -- The time interval at which the cache history is condensed for the following caches: TbeInferenceAgent, TbeObjectTable, TbeBackingStore, TbeNodeEvents, TbeAgentEvents, TbeNodeConcepts, and TbeClusterSummary. The default is 60 seconds.
- **Condense Raw Time** -- The time span of raw data kept in the cache history table for the following caches: TbeInferenceAgent, TbeObjectTable, TbeBackingStore, TbeNodeEvents, TbeAgentEvents, TbeNodeConcepts, and TbeClusterSummary. The default is 1200 seconds.

- **Compaction Rules** -- This field defines the rules used to condense your historical data in the database for the following caches: TbeInferenceAgent, TbeObjectTable, TbeBackingStore, TbeNodeEvents, TbeAgentEvents, TbeNodeConcepts, and TbeClusterSummary. By default, the columns kept in history will be aggregated by averaging rows with the following rule 1h -;1d 5m;2w 15m, which means the data from 1 hour will not be aggregated (1h - rule), the data over a period of 1 day will be aggregated every 5 minutes (1d 5m rule), and the data over a period of 2 weeks old will be aggregated every 15 minutes (2w 15m rule).

1. Navigate to the **Solution Package Configuration** > **TIBCO BusinessEvents** > **DATA STORAGE** tab.

2. In the **Compaction** region, click the **Condense Interval**, **Condense Raw Time**, and **Compaction Rules** fields and specify the desired settings.

   Note: When you click in the **Compaction Rules** field, the **Copy default text to clipboard** link appears, which allows you copy the default text (that appears in the field) and paste it into the field. This allows you to easily edit the string rather than creating the string from scratch.

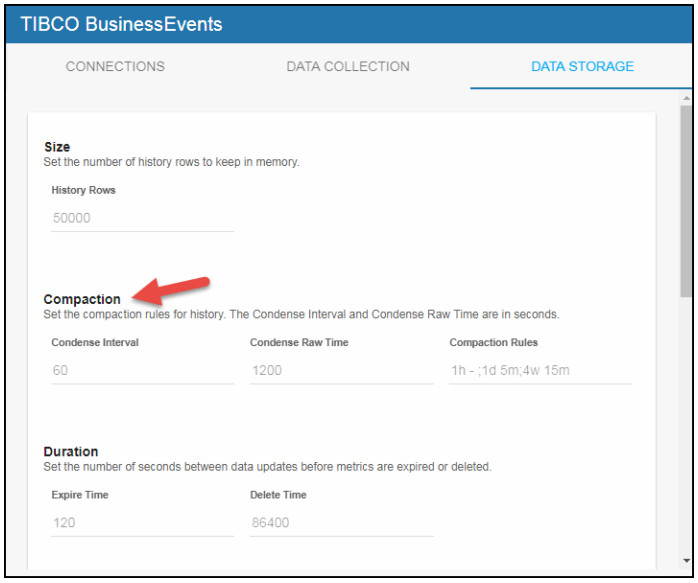

Defining Expiration and Deletion Duration for BEMON Metrics

The data for each metric is stored in a specific cache and, when the data is not updated in a certain period of time, that data will either be marked as expired or, if it has been an extended period of time, it will be deleted from the cache altogether. By default, metric data will be set to expired when the data in the cache has not been updated within 45 seconds. Also, by default, if the data has not been updated in the cache within 3600 seconds, it will be removed from the cache.

The caches impacted by the **Expire Time** and **Delete Time** properties are: TbeInferenceAgent, TbeClusterNode, TbeRtcTxnManagerReport, TbeObjectTable, TbeDbConnectionPool, TbeAvailability, TbeBackingStore, TbeNodeEvents, TbeAgentEvents, TbeNodeConcepts, TbeDestinationStatus, TbeChannelStatus, and TbeClusterSummary. To modify these defaults:

1. Navigate to **Solution Package Configuration** > **TIBCO BusinessEvents** > **DATA STORAGE** tab.

2. In the **Duration** region, click the **Expire Time** and **Delete Time** fields and specify the desired settings.

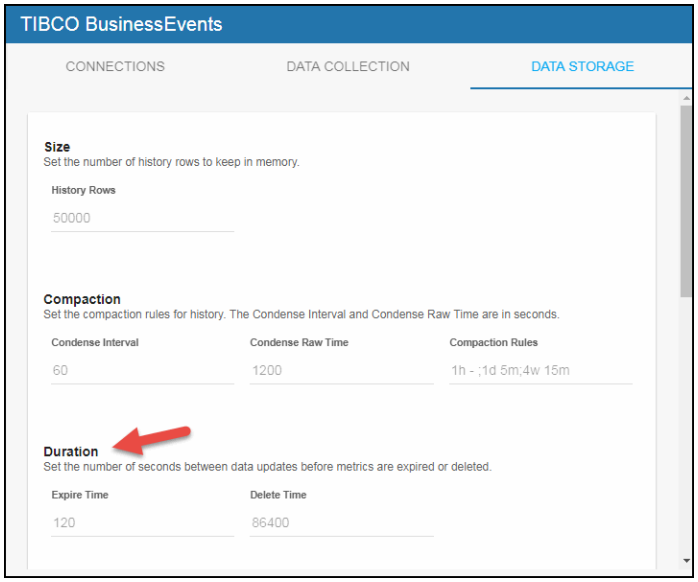

Enabling/Disabling Storage of BEMON Historical Data

The **History Storage** section allows you to select which metrics you want the Historian to store in the history database. By default, all historical data (in the TbeAgentEvents, TbeAvailability, TbeBackingStore, TbeClusterSummary, TbeInferenceAgent, TbeNodeConcepts, TbeNodeEvents, and TbeObjectTable caches) is saved to the database. To disable the collection of historical data, perform the following steps:

1. Navigate to **Solution Package Configuration** > **TIBCO BusinessEvents** > **DATA STORAGE** tab.

2. In the **History Storage** region, (de)select the toggles for the various metrics that you (do not) want to collect. Blue is enabled, gray is disabled.

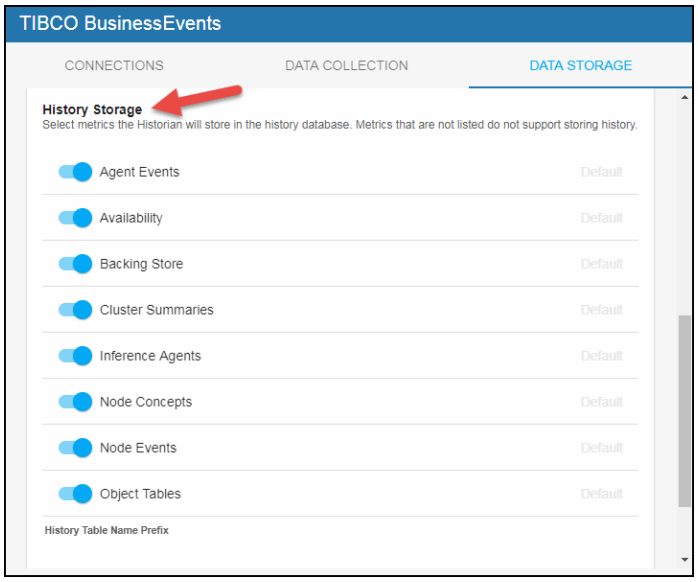

Defining a Prefix for All History Table Names for BEMON Metrics

The **History Table Name Prefix** field allows you to define a prefix that will be added to the database table names so that the Monitor can differentiate history data between data servers when you have multiple data servers with corresponding Historians using the same solution package(s) and database. In this case, each Historian needs to save to a different table, otherwise the corresponding data server will load metrics from both Historians on startup. Once you have defined the **History Table Name Prefix**, you will need to create the corresponding tables in your database as follows:

- Locate the .sql template for your database under **TIB_rtview-bw/rtvapm/tbemon/dbconfig** and make a copy of template.
- Add the value you entered for the **History Table Name Prefix** to the beginning of all table names in the copied .sql template.
- Use the copied .sql template to create the tables in your database.

To add the prefix:

1. Navigate to **Solution Package Configuration** > **TIBCO BusinessEvents** > **DATA STORAGE** tab.
2. Click on the **History Table Name Prefix** field and enter the desired prefix name.

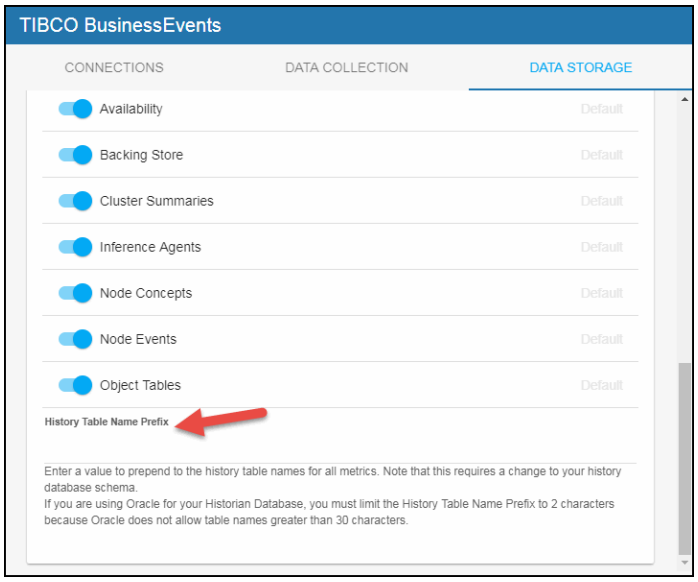

Configuring Ports

This section describes how to configure the ports for the RTView Servers. This step is required if you need to modify port settings or deploy Java processes on different hosts. Otherwise, this step is optional.

Java Processes

There are several Java processes included with the Monitor that are used during browser deployment. By default, it is assumed that these Java processes run on one host and that no configuration changes are needed. However, if these processes are distributed across several hosts, or if the default port definitions for these processes need to be modified, then configuration file settings must also be modified to allow all Monitor components to communicate with each other.

Java Process	Description	Default Port(s)
RTView Data Server	Gathers performance metrics.	Default Port= **3278** Default JMX Port = **3268**
RTView Historian	Retrieves data from the RTView Data Server and archives metric history to a database.	Default JMX Port= **3267**
RTView Display Server	Collects the data and generates the displays that the Application Server uses to produce the web pages.	Default Port= **3279** Default JMX Port = **3269**
Eclipse Jetty	Hosts the RTView Servlets.	Default Port= **3270**

To modify port settings or deploy Java processes on different hosts (rather than on a single host):

1. Navigate to the RTView Configuration Application > **(Project Name)** > **Server Configuration** > **General** > **GENERAL** tab.

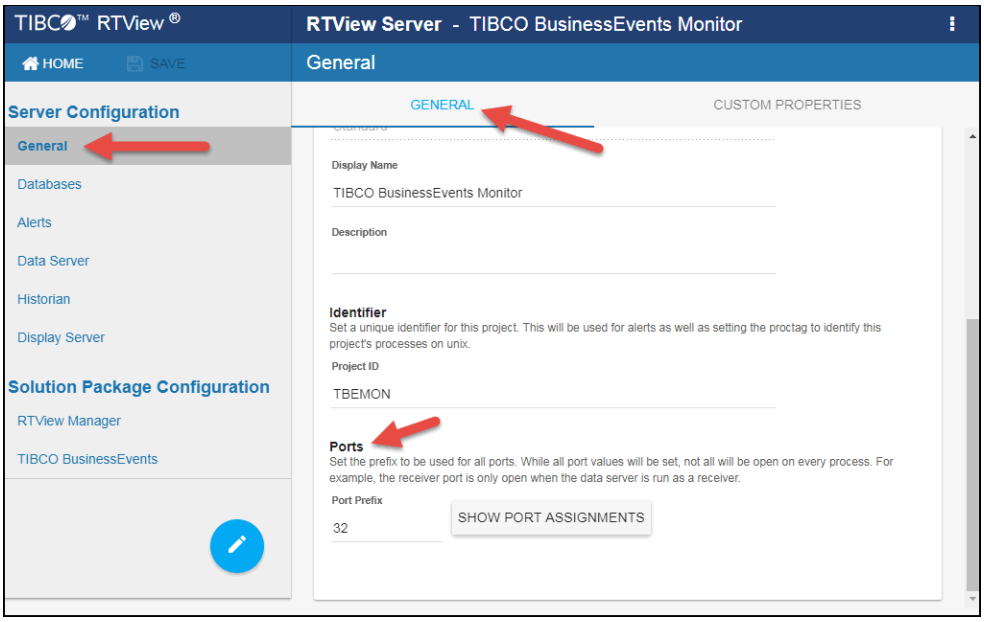

2. In the **Ports** region, click the **Port Prefix** field and specify the port prefix that you want to use. You can click the **SHOW PORT ASSIGNEMENTS** button to view the port numbers that are created using **Port Prefix** you specified.

3. Click **SAVE** in the RTView Configuration Application and restart the dataserver.

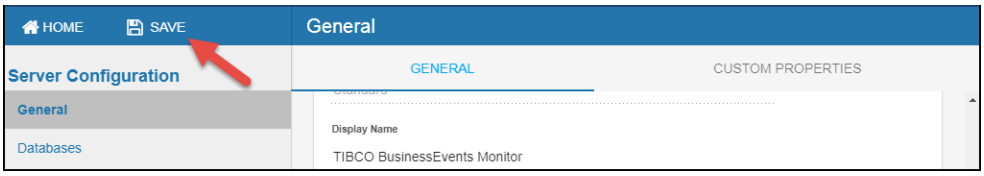

4. Edit the **update_wars** (**.bat** or **.sh**) file (in the **TIB_rtview-be/projects/rtview-server** directory) and change the port prefix for all ports to the prefix specified in Step 2.

5. Rebuild the war files for your application server by executing the following scripts (in the **TIB_rtview-be/projects/rtview-server** directory):

   Windows:

   update_wars.bat

   UNIX:

   ./update_wars.sh

6. Deploy your war files to your application server. For example:

   If you are using Tomcat, copy the **bemon.war** file located in the **TIB_rtview-be/projects/rtview-server** directory to the Tomcat **webapps** directory.

Configuring the Database

The Monitor is delivered with a default memory resident HSQLDB database, which is suitable for evaluation purposes. However, in production deployments, we recommend that you deploy one of our supported databases. For details, see the TIBCO© RTView© Standard Monitor User's Guide.

This section describes how to setup an alternate (and supported) database.

Database Requirements

The Monitor requires two database connections that provide access to the following information:

- **Alert Settings**

  The ALERTDEFS database contains alert administration and alert auditing information. The values in the database are used by the alert engine at runtime. If this database is not available, the Self-Service Alerts Framework under which alerts are executed will not work correctly.

- **Historical Data**

  The RTVHISTORY database contains the historical monitoring data to track system behavior for future analysis, and to show historical data in displays.

To Configure the Monitor Database:

You configure the database by defining database configurations in the RTView Configuration Application. You will also copy portions of the **database.properties** template file (located in the **common\dbconfig** directory) into the RTView Configuration Application.

1. Install a database engine of your choice. Supported database engines are Oracle, Microsoft SQL Server, MySQL, and DB2.

   NOTE: The default page size of DB2 is 4k. It is required that you create a DB2 database with a page size of 8k. Otherwise, table indexes will not work.

2. Open the **database.properties** template file, which is located in the **common\dbconfig** directory, and find the line that corresponds to your supported database from the "Define the ALERTDEFS DB" section.

3. Navigate to the RTView Configuration Application > **(Project Name)** > **Server Configuration** > **Databases** > **CONNECTIONS** tab and click the Edit icon in the **Alert Threshold Database Connection** region.

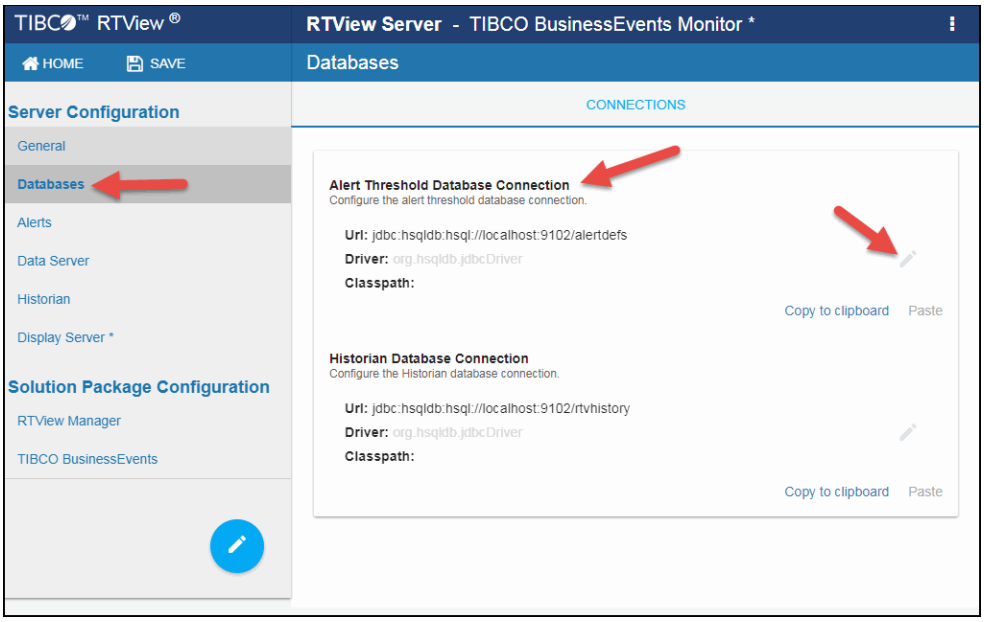

The **Edit Connection** dialog displays.

4. Enter the information from Step 2 into the **Edit Connection** dialog and click **Save**.

URL - Enter the full database URL to use when connecting to this database using the specified JDBC driver.

Driver - Enter the fully qualified name of the JDBC driver class to use when connecting to this database.

Classpath - Enter the location of the jar where the JDBC driver resides in your environment.

Username - Enter the username to enter into this database when making a connection.

Password - Enter the password to enter into this database when making a connection. If there is no password, use "**-**".

Run Queries Concurrently - Select this check box to run database queries concurrently.

5. Go back to the **database.properties** template file, which is located in the **common\dbconfig** directory, and find the line that corresponds to your supported database from the "Define the RTVHISTORY DB" section.

6. Navigate to the RTView Configuration Application > **(Project Name)** > **Server Configuration** > **Databases > CONNECTIONS** and then click the Edit icon in the **Historian Database Connection** region.

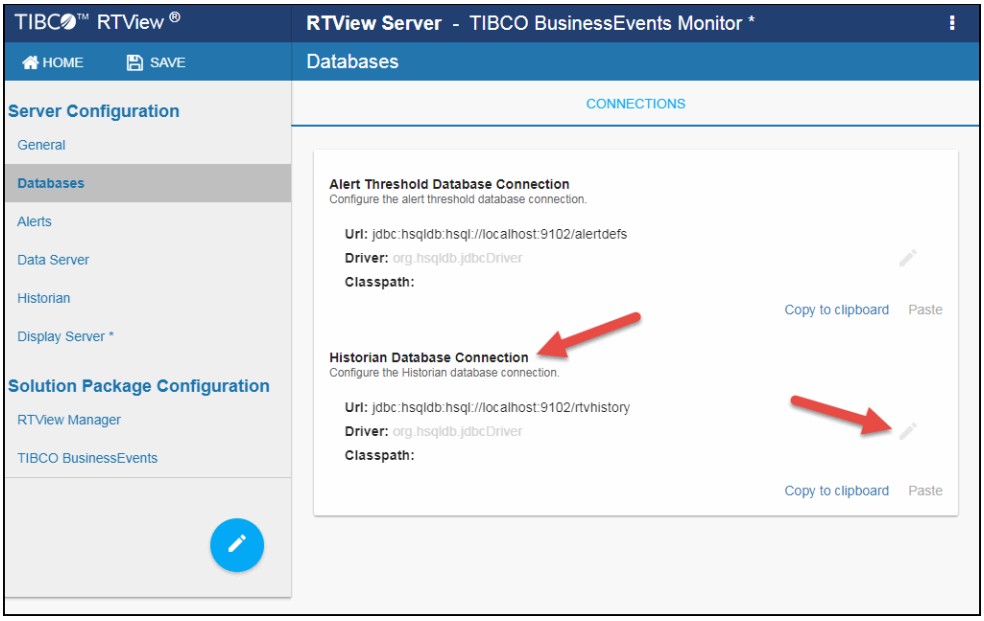

The **Edit Connection** dialog displays.

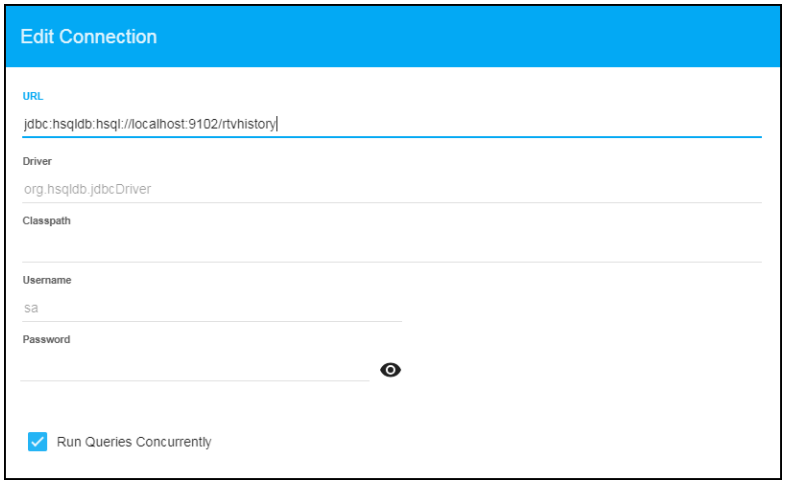

7. Enter the information from Step 5 into the **Edit Connection** dialog and click **Save**.

URL - Enter the full database URL to use when connecting to this database using the specified JDBC driver.

Driver - Enter the fully qualified name of the JDBC driver class to use when connecting to this database.

Classpath - Enter the location of the jar where the JDBC driver resides in your environment.

Username - Enter the username to enter into this database when making a connection.

Password - Enter the password to enter into this database when making a connection. If there is no password, use "**-**".

Run Queries Concurrently - Select this check box to run database queries concurrently.

8. Click **Save** in the RTView Configuration Application and restart your data server.

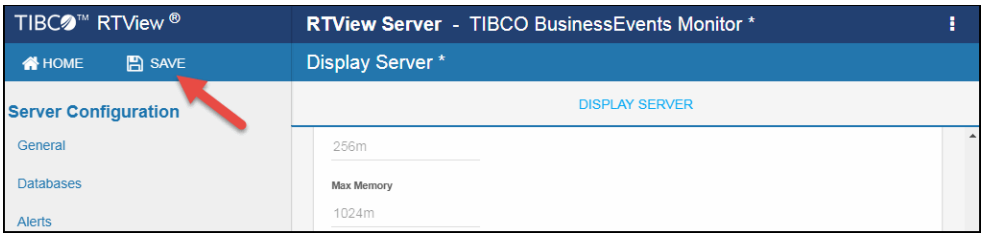

9. Manually create database tables. If your configured database user has table creation permissions, then you only need to create the Alerts tables. If your configured database user does not have table creation permission, then you must create both the Alert tables and the History tables.

To create tables for your database, use the **.sql** template files provided for each supported database platform, which is located in the **dbconfig** directory of the **common**, **bemon**, and **rtvmgr** directories:

- **Alerts**

rtvapm/common/dbconfig/create_common_alertdefs_tables_<db>.sql

- **History**

rtvapm/bemon/dbconfig/create_bemon_history_tables_<db>.sql

rtvapm/rtvmgr/dbconfig/create_rtvmgr_history_tables_<db>.sql

where **<db>** ={**db2**, **mysql**, **oracle**, **sqlserver**}

NOTE: The standard SQL syntax is provided for each database, but requirements can vary depending on database configuration. If you require assistance, consult with your database administrator.

The most effective method to load the **.sql** files to create the database tables depends on your database and how the database is configured. Some possible mechanisms are:

- **Interactive SQL Tool**
  Some database applications provide an interface where you can directly type SQL commands. Copy/paste the contents of the appropriate **.sql** file into this tool.
- **Import Interface**
  Some database applications allow you to specify a **.sql** file containing SQL commands. You can use the **.sql** file for this purpose.

Before loading the **.sql** file, you should create the database and declare the database name in the command line of your SQL client. For example, on MySQL 5.5 Command Line Client, to create the tables for the Alert Settings you should first create the database:

create database myDBName;

before loading the **.sql** file:

mysql -u myusername -mypassword myDBName < create_common_alertdefs_tables_mysql.sql;

If you need to manually create the Historical Data tables, repeat the same process. In some cases it might also be necessary to split each of the table creation statements in the **.sql** file into individual files.

Third Party Application

If your database does not have either of the two above capabilities, a third party tool can be used to enter SQL commands or import **.sql** files. Third party tools are available for connecting to a variety of databases (RazorSQL, SQLMaestro, Toad, for example).

You have finished configuring the databases.

Enable Monitoring of TIBCO APIX Engines

TIBCO API Exchange Gateway (APIX) uses engines similar to TIBCO BusinessEvents engines, and you can configure TIBCO BusinessEvents Monitor to monitor TIBCO APIX engines by performing the following:

- Include two new files via **cache.config** properties:

```
tbe_cache_source_be5.1_asg_cacheagent.rtv
tbe_cache_source_be5.1_asg_inferenceagent.rtv
```

Note: You will need separate connections for cache and inference to a single engine process.

- Include properties like the following in your runtime properties file (where **asgCacheConn** and **asgInferConn** are arbitrary but distinct names. For example: "asgcache_4456, asginter_4456”):

```
#collector.sl.rtview.jmx.jmxconn=<asgCacheConn> <hostname> <port> URL:- - -
'false'
#collector.sl.rtview.jmx.jmxconn=<asgInferConn> <hostname> <port> URL:- - -
'false'
#collector.sl.rtview.cache.config= tbe_cache_source_be5.1_asg_cacheagent.rtv
$tbe_conn:<asgCacheConn> $tbe_cluster:<clusterName>
#collector.sl.rtview.cache.config=tbe_cache_source_be5.1_asg_inferenceagent.rtv
$tbe_conn:<asgInferConn> $tbe_cluster:<clusterName>
```

Enable Collection of Tomcat Historical Data

By default, collection of Tomcat historical data is disabled. To enable collection of Tomcat history:

1. Navigate to RTView Configuration Application > (Project) > **Solution Package Configuration** > **RTView Manager** > **DATA STORAGE** tab.
2. Select the **Tomcat Global Requests**, **Tomcat Webmodule Statics**, and/or **Tomcat Webmodule Totals** to enable collecting Tomcat historical data.

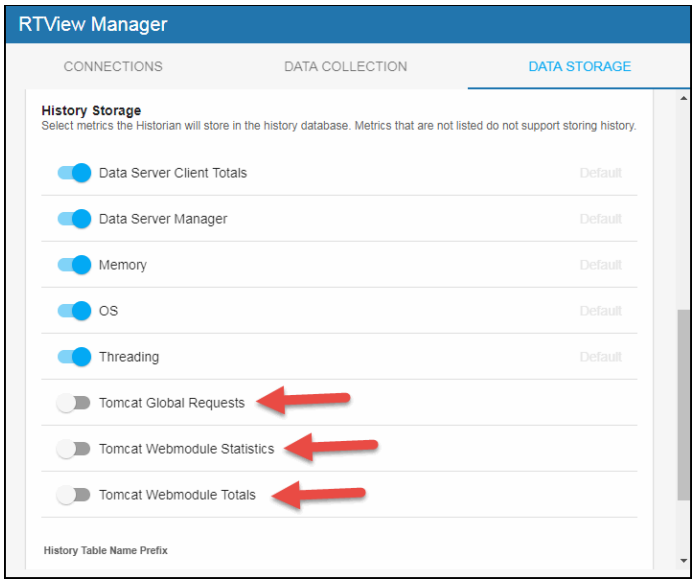

Configure Alert Notification

This section describes how to configure alerts to execute an automated action (such as sending an email alert). To setup alert notification you select the event you want to notify on and then select the action to execute.

You set alerts to execute notifications based on the following events:

- when a new alert is created
- the first time the **Severity** level on an alert changes
- when an alert is cleared
- periodically renotify unacknowledged alerts

By default, a **.bat** script is executed for new alerts and on the first severity change for an alert. The script, by default, is not configured to execute an automated action. However, you can uncomment a line in the script that prints alert data to standard output. Or, you can modify the script to execute an automated action (such as sending an email alert). The following is a sample output from the alert command script:

----- Alert command script executed: DOMAINNAME=MYMON-1, ALERTNAME=someAlert, ALERTINDEXX=alertIndex1~alertIndex2, ALERTID=1075, ALERTSEVERITY=2, ALERTTEXT=High Alert Limit exceeded current value: 100.0 limit: 80.0 #####

To configure Alert Notification:

1. Open the RTView Configuration Application, select **Alerts** (in the navigation tree) and then the **Alerts** tab.

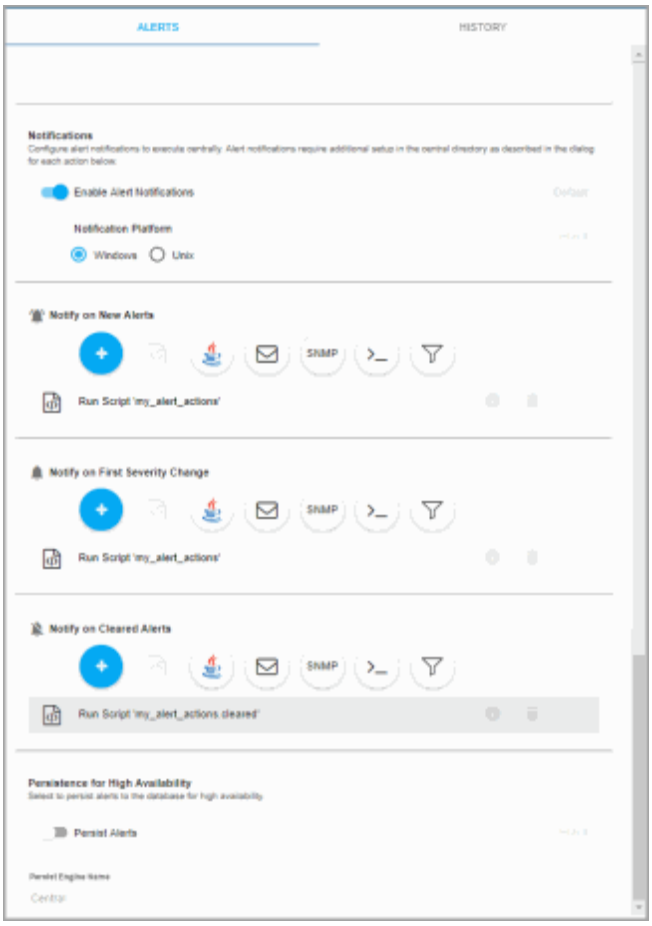

2. Toggle on **Enable Alert Notifications** and select the **Notification Platform** type (**Windows or Unix**).

3. Select an alert event that you want to notify on by clicking next to the option.

Alert Event Options

- **Notify on New Alerts**:A notification is executed every time a new alert is created.
- **Notify on First Severity Change**: A notification is executed the first time the **Severity** changes for each alert.
- **Notify on Cleared Alerts**: A notification is executed every time an alert is cleared.
- **Periodically Renotify on Unacknowledged Alerts**: Enter the **Renotification Interval** (number of seconds). A notification is executed for each unacknowledged alert per the interval you specify here. If the Renotification Interval is greater than **0** and no actions are defined, the **New Alerts** action will be used for renotifications.

4. Select the alert action(s) you want to execute.

Alert Action Options

Run a Script

Execute Java Code

 Add Email Notification

Send SNMP Trap

Run Command String

Conditional Filter

You can choose multiple actions.

5. Click to close the dialog and (in title bar) to save your changes.

6. Some alert notification actions require additional setup as described in the dialog for each action. See the descriptions of each action below for details on the dialogs and additional setup for each action.

7. Click to apply changes.

Run a Script

This alert notification action executes the following script in the **projects/rtview-server** directory for RTViewCentral and in the **projects/rtview-manager** directory for RTView Manager:

- **my_alert_actions.bat/sh** – New and First Severity Change
- **my_alert_actions.cleared.bat/sh** – Cleared
- **my_alert_actions.renotify.bat/sh** – Periodically Renotify

This action can only be added once per notification type. In addition to selecting this action in the Configuration Application, you must also modify the appropriate script to execute the actions for your notification. This script has access to the following fields from the alert: **Alert Name**, **Alert Index**, **ID**, **Alert Text** and **Severity**.

Return to Alert Event Options.

Execute Java Code

This alert notification action allows you to implement your alert notification actions using Java code. It executes the **my_alert_notification.$domainName.$alertNotifyType.$alertNotifyCol** command in your Custom Command Handler and passes the row from the alert table that corresponds to the alert.

This action can only be added once per notification type. In addition to selecting this action the Configuration Application you must also modify the custom command handler to execute the actions for your notification. A sample custom command handler is included under **projects/custom**. It prints the alert notification to the console. You will modify this command handler to implement your own notification actions.

Make the following entries:

- **Custom Command Handler Class Name**: Enter the fully qualified name of the Custom Command Handler class. This defaults to the sample Custom Command Handler in the **projects/custom** directory.
- **Custom Command Handler Jar**: Enter the path and name of the jar containing the Custom Command Handler class. The path may be absolute or relative to the location

of data server. This defaults to the sample Custom Command Handler in the **projects/custom** directory.

Note that if you can only have one custom command handler per Data Server, so changing these settings for one notification event will change them for the rest of the notification events.

Customizing the Custom Command Handler

The source for the Custom Command handler is provided in the **RtvApmCommandHandler.java** file, located in the **RTViewEnterpriseMonitor\projects\custom\src\com\sl\rtvapm\custom** directory. By default, the handler prints the alert data to standard output. To change this behavior perform the following steps:

1. Open the **RtvApmCommandHandler.java** file.
2. Modify the **OutputAlertString** method as needed. You can replace this method with your own if you modify the **invokeCommand** method to call it, and your method accepts the same arguments as **OutputAlertString**.
3. Save the **RtvApmCommandHandler.java** file.
4. Compile **RtvApmCommandHandler.java** and rebuild **rtvapm_custom.jar** using the supplied script (**make_all.bat** or **make_all.sh**) in **projects\custom\src** directory.

Return to Alert Event Options.

Add Email Notification

This alert notification action sends an email. This action can be added multiple times per notification type. No additional setup is required beyond filling in the **Add Email Notification** dialog in the Configuration Application.

Make the following entries:

- **SMTP Host**: The SMTP host address. This is required. Consult your administrator.
- **SMTP Port**: The SMTP port number. This is required. Consult your administrator.
- **From**: The email address from which to send the email. This is required.
- **To**: The email address to which to send the email. This is required and may contain multiple entries.
- **Subject**: The subject for the email. This is required. You can include the value from any column in the alert table in your subject. Click **Insert $alert<Value>** and select one or more applicable alert value(s).
- **Body**: The body of the email. This is optional. Click **Insert $alert<Value>** and select one or more applicable alert value(s).
- **User**: The user name for the account from which you are sending the email. This is optional.
- **Password**: The password for the account from which you are sending the email. This is optional.

Return to Alert Event Options.

Send SNMP Trap

This alert notification action sends an SNMP Trap as described in **rtvapm/common/lib/SL-RTVIEW-EM-MIB.txt**. This action can be added multiple times per notification type. No additional setup is

required beyond filling in the **Add SNMP Trap Notification** dialog in the Configuration Application.

Make the following entries:

- **Trap Type**: Select the SNMP version of the trap. This is required.
- **Destination Address**: The system name or IP address of the receiving system. This is required.
- **Destination Port**: The UDP port on the receiving system. This is required.
- **Community Name**: (This field is visible when **Trap Type v2/v3** is selected.) The SNMP v2 Community Name string. This is required.

Return to Alert Event Options.

Run Command String

This alert notification action executes a specified command. This action can be added multiple times per notification type. Make the following entry:

Command String: Enter the command string for any command supported by RTView. To enter a command string, you must know the correct syntax for the command. Contact Technical Support for assistance on syntax. You can include the value from any column in the alert table using the syntax in the Show More link at the bottom of the dialog.

Return to Alert Event Options.

Conditional Filter

This alert notification action alert allows you to execute different actions for different alerts based on information in the alert. For example, you can configure EMS alerts to send emails to your EMS team and Solace alerts to send emails to your Solace team. This action can be added multiple times per notification type.

To create a condition, make the following entries:

- **Alert Field**: Select an alert field: **Alert Name**, **Alert Index**, **Category, CI Name, Owner, Package, Primary Service** or **Severity**. This is required.
- **Operator**: Select one - **EQUALS**, **DOES NOT EQUAL**, **STARTS WITH**, **ENDS WITH** or **CONTAINS**. This is required.
- **Value**: Enter the value to which to compare the Alert Field. Cannot contain wildcard characters. This is required.
- **Action(s)**: Select one or more actions to execute when this condition is met - Run a Script , Execute Java Code , Send SNMP Trap , Add Email Notification or Run Command String.

Return to Alert Event Options.

Configure High Availability

High Availability (HA) mitigates single point of failure within TIBCO Business Events Monitor by providing a means of defining redundant system components, together with failover capability, for users of those components.

When using HA, components are designated **PRIMARY** and **BACKUP**. If the **PRIMARY** component fails, failover occurs to the **BACKUP** component. If the **PRIMARY** component is subsequently restarted, the **BACKUP** component allows the newly restarted component to take the primary role and return to its backup role.

This section contains the following:

- Overview of High Availability Architecture
- Requirements for Configuring High Availability
- Steps for Configuring High Availability
- Verifying the High Availability Configuration

Overview of High Availability Architecture

Data Server High Availability

The primary and backup data servers connect to each other via socket. If the primary data server stops, then the backup server takes over. If the primary then comes back online, then the primary takes over again and the backup returns to standby mode. The data client connections will move between the two servers accordingly.

NOTE: Be aware that data clients can connect to the standby server using a non-fault tolerant URL and still get data because of a proxy feature where the standby server forwards data requests to the primary server. This can be confusing when you use the HTML Cache Viewer (**http://localhost:3270/common**) on the standby server to view cache contents because it looks like the standby server caches are updating, but you are really viewing the data in the primary server and not in the standby server.

Display Server High Availability (Classic User Interface)

In display server deployments, the primary display server and backup display server do not connect to each other. The rtvdisplay servlet is configured to connect first to the primary and, if that fails, it tries to connect to the backup. At any point, if the one it is connected to becomes unavailable, then it will try to connect to the other. You can configure whether to have the rtvdisplay server connect back to the primary server when it comes back online or stay connected to the backup server until it goes offline.

HTML User Interface High Availability

The HMTL UI client connects to the data server via an HA configured rtvquery servlet.

Historian High Availability

The primary and backup historian connect to each other via socket. If the primary historian stops, then the backup takes over. If the primary historian comes back online, then the primary takes over again and the backup returns to standby mode. Only the active historian writes to the database.

The historian is a data client of the data server and connects to it via a fault tolerant URL (socket only), which means that the data servers and historians can fail over separately or together.

Requirements for Configuring High Availability

The following are minimum requirements for High Availability:

- Two host machines, one for the primary host and one for the backup host.
- Both hosts must be configured such that the RTView processes on each host can connect to each other via socket.
- Both hosts must be able to access:
  - the same data connections
  - the same historian database
  - the alert threshold database
- The RTView processes on both hosts must be able to run against identical properties files. In the case where drivers or other third party jars are located in different directories on the two hosts, create a directory in the same location in each host, copy the jar files into and reference that directory in your properties.
- Tomcat or other Application Server
  - The HTML UI and rtv servlets must be deployed on an application server other than the internal Jetty server. Note that this requires extra configuration of the servlet **.war** files in the application server.

Steps for Configuring High Availability

To Configure High Availability:

1. On both the primary and backup hosts, define the following environment variables:
   - **PRIMARYHOST** - the IP Address or hostname of the host running the primary servers (for example, **set PRIMARYHOST=MyHost**).
   - **BACKUPHOST** - the IP Address or hostname of the host running the backup servers (for example, **set BACKUPHOST=OtherHost**).
2. Install the Monitor on both the primary host and the backup host.
3. Configure your servlets to be HA and deploy them to your application server:
   - **cd projects\rtview-server**
   - In a text editor, open **update_wars**(**.bat** or **.sh**) and fill in the values for **HOST**, **HA_HOST**, **HA_DISPLAYHOST**, and **HA_FAILBACK** as described in the script.
   - Run the **update_wars**(**.sh** or **.bat**) script.
   - Copy the generated war files to the **webapps** directory of your application server.
4. To run High Availability, you must run the following from the command line:

Windows

- From the command line on the primary host, **cd** to <top level directory> and type **start_server -haprimary**.
- From the command line on the backup host, **cd** to <top level directory> and type **start_server -habackup**.

Unix

- From the command line on the primary host, **cd** to <top level directory>and type **start_server.sh -haprimary**.
- From the command line on the backup host, **cd** to <top level directory> and type **start_server.sh -habackup**.

5. Configure the Monitor on the primary host using the RTView Configuration Application (see Quick Start and RTView Configuration Application for more information on configuring the Monitor). Make sure to configure data collection, configure server options and databases, and enable alert persistence.

Note that the RTView Configuration Application must be able to connect both the primary and backup servers in order to enable editing. The same properties are saved to both servers. The **RESTART SERVERS** button (in the RTView Configuration Application) restarts both the primary and backup servers at the same time. If you want to stagger the restarts, use the scripts under **TIB_rtview-bw** to stop and then start your servers after making changes in the RTView Configuration Application.

Note: Jetty does not have to be disabled, but data clients will not be able to make high availability connections to the data server using the Jetty URL. However, the Jetty URL can still be used to configure the application.

Verifying the High Availability Configuration

Verify failover and failback configurations by looking for the following in the log files.

Note: If the PRIMARYHOST and/or BACKUPHOST environment variable(s) is/are not set, you will get the following error in the log files and HA will be disabled:

```
ERROR: Disabling HA because the PRIMARYHOST and/or BACKUPHOST environment
variable is not set.
```

Primary Data Server Log File

```
startup

[rtview] Starting as primary HA data server accessible via
//primaryhostname:3278,//backuphostname:3278

[rtview] DataServerHA: connected to backuphostname:3278

[rtview] DataServerHA: run as primary server, backuphostname:3278 has lower
priority than this server

[rtview] leaving standby mode
```

Backup Data Server Log File

```
startup
```

```
[rtview] Starting as backup HA data server accessible via
//primaryhostname:3278,//backuphostname:3278

rtview] entering standby mode

after failover (primary data server exits)

[rtview] DataServerHA: error receiving message: java.net.SocketException:
Connection reset (primaryhostname:3278)

[rtview] DataServerHA: becoming primary server, lost connection to primary server
primaryhostname:3278

[rtview] leaving standby mode

after failback (primary data server comes back up)

[rtview] DataServerHA: resigning as primary server, got standby directive from
other server primaryhostname:3278

[rtview] connected to primaryhostname:3278

[rtview] entering standby mode
```

Primary Historian Log File

```
[rtview] Starting as primary HA historian paired with backup historian at
<backuphostname>:3222

[rtview] ServerGroup: status of member <backuphostname>:3222: primary, priority= 1,
started=Wed Nov 14 12:56:01 PST 2018

[rtview] ServerGroup: primary server = local

[rtview] ServerGroup: becoming primary server
```

Backup Historian Log File

```
[rtview] Starting as backaup HA histoiran paired with primary historian at
<primaryhostname>:3222

[rtview] ServerGroup: status of member <primaryhostname>:3222: primary, priority= ,
started=Wed Nov 14 12:56:01 PST 2018

[rtview] ServerGroup: primary server = <primaryhostname>:3222

after failover (primary historian exits):

[rtview] error receiving message: java.io.EOFException (primaryhostname:3222)

[rtview] ServerGroup: disconnected from primaryhostname:3222

[rtview] ServerGroup: primary server = local

after failback (primary historian starts back up):

[rtview] ServerGroup: status of member primaryhostname:3222: primary, priority= 2,
started= Tue Nov 20 09:12:43 PST 2018

[rtview] ServerGroup: connected to primaryhostname:3222

[rtview] ServerGroup: primary server = primaryhostname:3222
```

Primary Display Server Log File

```
2018-11-19 14:08:09,366 INFO main - [rtview] Starting as primary HA display server
paired with backup display server on <backuphostname>
```

Backup Display Server Log File

```
2018-11-19 14:08:09,366 INFO main - [rtview] Starting as backup HA display server
paired with primary display server on <primaryhostname>
```

CHAPTER 4 Deployment

This section describes how to deploy the Monitor components. This section includes:

- "Overview" on page 42
- "Web Application Deployment" on page 42
- "RTView Server Components as Windows Services" on page 44
- "Troubleshooting" on page 47
- "Sender/Receiver: Distributing the Load of Data Collection" on page 48

Overview

The Monitor can be deployed as a stand-alone desktop client or as a web application that runs in a browser. Evaluation environments can use the provided HSQLDB database. Production environments require a supported JDBC- or ODBC-enabled relational database to store historical information. Supported databases are MySql, Oracle, SqlServer and DB2.

The RTView Historian and RTView Data Server are typically deployed on the same host. However, these processes can optionally be configured on separate hosts. Doing so can increase performance in deployments that need to support many end users or systems with large TIBCO servers.

- "Web Application Deployment" on page 42

If you choose the browser option, clients need only a browser and Adobe Flash installed. The RTView Display Server, RTView Data Server, RTView Historian and Application Server are typically installed on the same host.

Web Application Deployment

This section describes how to deploy the Monitor as a web application. You start the Monitor using the **start_server** script (and stop the Monitor using the **stop_server** script). For web application deployments the following processes are started: the RTView Data Server, Historian, and Display Server, as well as the database and an application server.

This section contains:

To deploy the Monitor as a web application:

- "Windows" on page 42
- "UNIX/Linux" on page 43

Windows

Note: You can skip Step 1 and Step 2 if you are using Eclipse Jetty, which is delivered with the Monitor, as your application server.

1. Copy the **.war** files, located in the **TIB_rtview-be\projects\rtview-server** directory, and deploy them to your Application Server.

2. Start your Application Server if using Tomcat or an application server other than Eclipse Jetty.

3. You can skip this step if you are using Eclipse Jetty. The RTView Configuration Application uses digest authentication for security, and only allows access to users with the "rtvadmin" role. In order to allow access to the RTView Configuration Application in your application server, you need to add a user with the "rtvadmin" role. For example, if using Tomcat, follow the instructions below. For other application servers, refer to their documentation for adding users.
   - Edit <installation directory>\conf\tomcat-user.xml
   - Add the following lines inside the tomcat-users tag:

```
<role rolename="rtvquery"/>
<user username="rtvquery" password="rtvadmin" roles="rtvquery"/>
```

4. Change directory (**cd**) to the **TIB_rtview-be** directory and start the Monitor applications by typing **start_server**.

NOTE: The **start_server.bat** command starts all the Monitor applications at once. Use the **stop_server.bat** script to stop Monitor applications.

5. Open a Web browser and access the following URL to open the Monitor:

   http://localhost:3270/rtview-tbemon (for the HTML displays)

   or

   http://localhost:3270/rtview-tbemon-classic (for the classic displays)

   If using your own application server:

   http://host:port/rtview-tbemon (or **rtview-tbemon-classic** for the classic displays)

   Where **host** is the IP or host name where your Application Server is running and **port** is the port used by your Application Server. The login display opens in the Web browser.

   Login. The default user name and password are:

   User Name: **admin**

   Password: **admin**

   The main Monitor display opens.

   See Quick Start for a more detailed example.

UNIX/Linux

> **Note:** You can skip Step 1 and Step 2 if you are using Eclipse Jetty, which is delivered with the Monitor, as your application server.

1. Copy the **.war** files, located in the **TIB_rtview-be\projects\rtview-server** directory, and deploy them to your Application Server.

   Note: You can skip this step if you are using Eclipse Jetty, which is delivered with the Monitor, as your application server.

2. Start your Application Server if using Tomcat or an application server other than Eclipse Jetty, which is delivered with the Monitor.

3. You can skip this step if you are using Eclipse Jetty. The RTView Configuration Application uses digest authentication for security, and only allows access to users with the "rtvadmin" role. In order to allow access to the RTView Configuration Application in your application server, you need to add a user with the "rtvadmin" role. For example, if using Tomcat, follow the instructions below. For other application servers, refer to their documentation for adding users.
   - Edit <installation directory>\conf\tomcat-user.xml
   - Add the following lines inside the tomcat-users tag:

```
<role rolename="rtvquery"/>
<user username="rtvquery" password="rtvadmin" roles="rtvquery"/>
```

4. Change directory (**cd**) to the **TIB_rtview-be** directory and start the Monitor applications by typing **start_server**.

NOTE: The **start_server.sh** command starts all the Monitor applications at once. Use the **stop_server.sh** script to stop Monitor applications.

5. Open a Web browser and access the following URL to open the Monitor:

   http://localhost:3270/rtview-tbemon (for the HTML displays)

   or

   http://localhost:3270/rtview-tbemon-classic (for the classic displays)

   If using your own application server:

   http://host:port/rtview-tbemon (or **rtview-tbemon-classic** for the classic displays)

   Where **host** is the IP or host name where your Application Server is running and **port** is the port used by your Application Server. The login display opens in the Web browser.

   Login. The default user name and password are:

   User Name: **admin**

   Password: **admin**

   The main Monitor display opens.

   See Quick Start for a more detailed example.

RTView Server Components as Windows Services

This section describes how to configure an RTView process (Data Server, Historian, Display Server) to run as a Windows service.

To Configure the Data Server, Historian or Display Server to run as a Windows Service

1. Navigate to the RTView Configuration Application > **(Project Name)** > **Server Configuration** > **General** > **Custom Properties** tab.

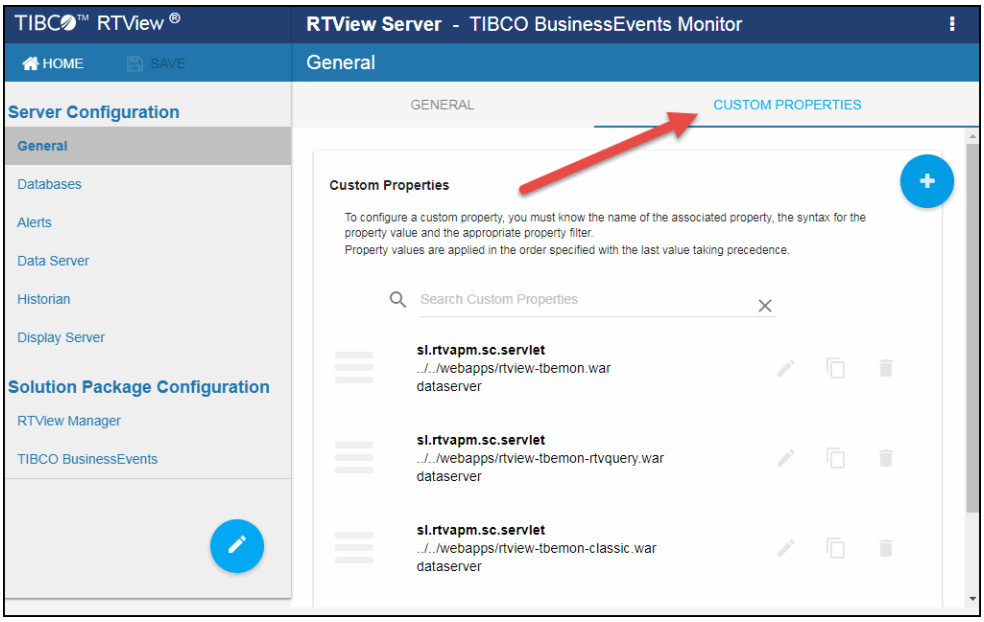

2. Click the icon.

The **Add Property** dialog displays.

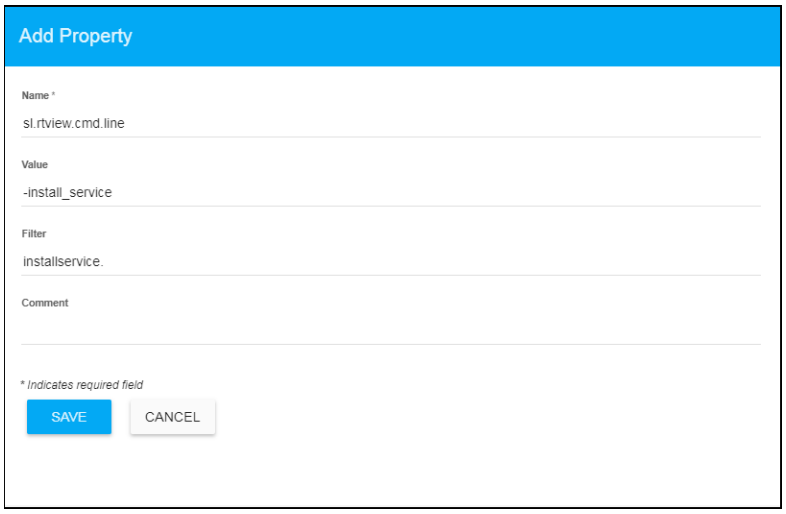

3. Define the values for each of the following properties and click **Save**:

Name: sl.rtview.cmd_line

Value: -install_service

Filter: installservice

Comment: (description of the filter)

Name: sl.rtview.cmd_line

Value: -dir:%RTVAPM_STARTUP%

Filter: installservice

Comment: (description of the filter)

Name: sl.rtview.cmd_line

Value: -uninstall_service

Filter: uninstallservice

Comment: (description of the filter)

Note: The environment variable %RTVAPM_STARTUP% is set by run script to the directory where the script was started.

4. For each Windows service you want to create, add the following property and replace ServiceName in the value and filter fields with a name you choose for the service:

Name: sl.rtview.cmd_line

Value: -service:ServiceName

Filter: ServiceName

For example, choose TBEMonData as the name for starting a Data Server as a Windows service and TBEMonDisp to indicate a name for starting a Display Server as a Windows service.

Name: sl.rtview.cmd_line

Value: -service:TBEMonData

Filter: TBEMonData

Name: sl.rtview.cmd_line

Value: -service:TBEMonDisp

Filter: TBEMonDisp

Note: Each service must have a unique name and the beginning of the property entered must match the name of the service.

5. Once all your properties have been added, click **SAVE**.

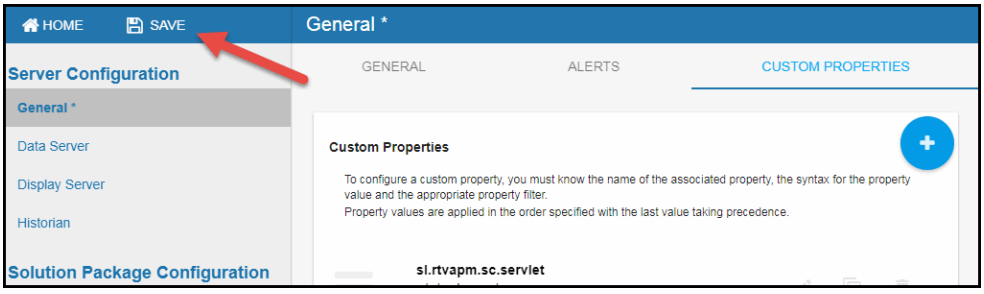

6. Restart the data server so that your changes take effect.

To install and run

7. Execute the following scripts to start the service:

NOTE: These scripts must be run in an initialized command window.

rundata -propfilter:installservice -propfilter:TBEMonData

rundisp -propfilter:installservice -propfilter:TBEMonDisp

To uninstall

8. Execute the following scripts to uninstall the services:

NOTE: These scripts must be run in an initialized command window.

rundisp -propfilter:uninstallservice -propfilter:TBEMonDisp

rundata -propfilter:uninstallservice -propfilter:TBEMonData

Troubleshooting

This section includes:

- "Log Files," {Default ¹ Font}next
- "JAVA_HOME" on page 47
- "Permissions" on page 47
- "Network/DNS" on page 47
- "Verify Data Received from Data Server" on page 48
- "Verify Port Assignments" on page 48

Log Files

When a Monitor component encounters an error, it outputs an error message to the console and/or to the corresponding log file. If you encounter issues, look for errors in the following log files:

- **dataserver.log**
- **displayserver.log**
- **historian.log**

which are located in the **TIB_rtview-be\projects\rtview-server\logs** directory.

Logging is enabled by default. If you encounter issues with log files, verify the **logs** directory exists in the **TIB_rtview-be\projects\rtview-server\logs** directory.

JAVA_HOME

If the terminal window closes after executing the **start_server** command, verify that JAVA_HOME is set correctly.

Permissions

If there are permissions-related errors in the response from the **start_server** command, check ownership of the directory structure.

Network/DNS

If any log file shows reference to an invalid URL, check your system's hosts file and confirm with your Network Administrator whether your access to the remote system is being blocked.

Verify Data Received from Data Server

If you encounter problems collecting data, restart the Data Server, start the Monitor, and go to the **Admin** tab and select **Architecture> RTView Cache Tables** in the navigation tree. Select **TBEMON-LOCAL** from the **Data Server** drop down list, and search for all caches that start with "TBEMON." Make sure these caches are populated (the number of **Rows** and **Columns** in the table should be greater than 0). If not, there might be a problem with the connection to the Data Server.

Verify Port Assignments

If the display server or Historian fail to connect to the Data Server or they receive no data, verify the ports are assigned correctly in your properties files and restart the Data Server.

Sender/Receiver: Distributing the Load of Data Collection

The standard method of collecting data involves one Data Server that sends the data to the Display Server and the Historian. For example:

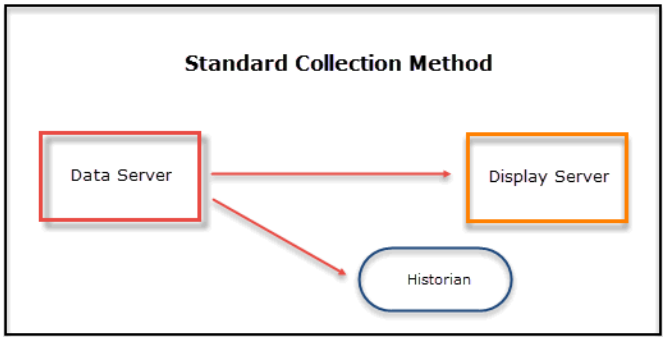

This method is optimized to deliver data efficiently when large tables and high data volumes are involved. There is, however, an alternative method of collecting data: the Sender/Receiver Data Collection Method. This collection method allows you to configure TIBCO Business Events Monitor so that you have a Data Server (Receiver) that collects data from one or more remote Senders. This type of configuration could be useful in the following scenarios:

1. When dividing the collection load across different machines is more efficient

   In the Sender/Receiver Data Collection Method, the Senders are configured as lightweight Data Servers without history being configured and whose primary purpose is to collect and aggregate data from their respective local servers that they then send to the full-featured Data Server (Receiver). The benefit of this type of configuration comes from balancing the load of the data collection. The Senders collect data exclusively from the servers in their network and send the data to the Receiver, which collects the data and sends it to the TIBCO Business Events Monitor Display Server, the Historian, and the Viewer. The following illustration provides one configuration example:

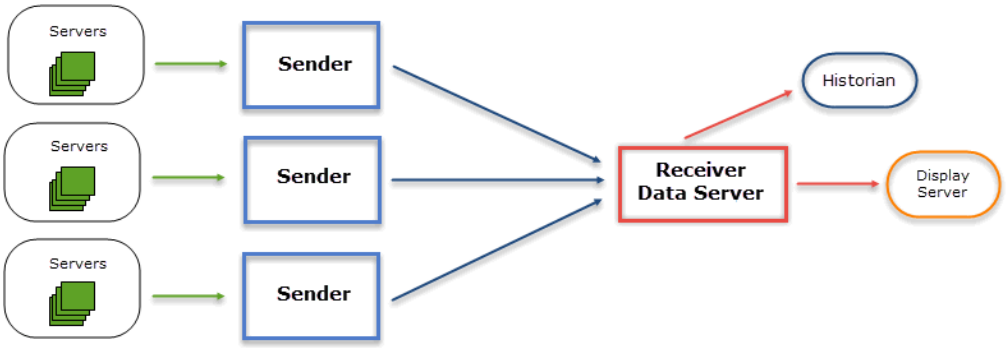

2. When firewall limitations prevent the Receiver Data Server from receiving data directly, Senders behind the firewall can be configured to send data to the Receiver

   In the Standard Collection Method, the client must specify the network address of the Data Server to which it wants to connect, which might not be allowed due to security restrictions. In these situations, the Sender/Receiver Collection Method could be considered since the Receiver does not need to know the network addresses of the Senders because it simply opens the port and passively receives data from any defined Sender.

Example

The following example contains Senders in London and Madrid that collect data from their associated servers and send the data to a Receiver Data Server in New York City. The Receiver takes the collected data from London and Madrid along with data collected from its own associated servers and sends it to the TIBCO Business Events Monitor displays.

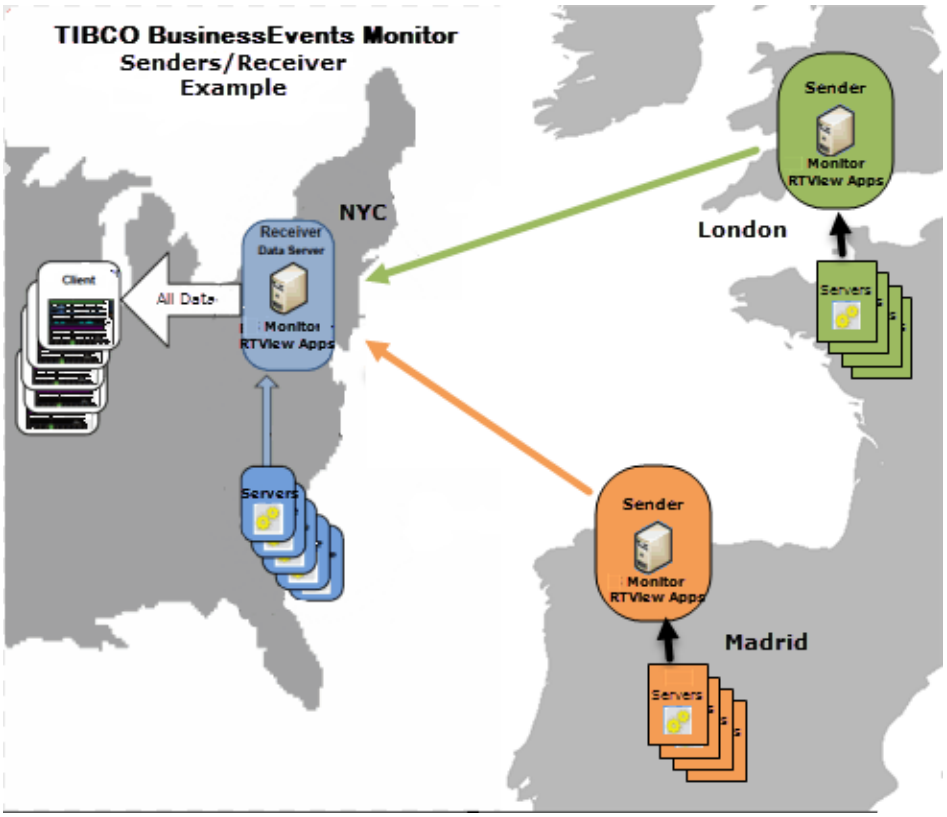

Receiver Data Server -- NYC	Sender -- London	Sender -- Madrid
• Automatically detects and gathers data from its local servers. • Receives data from London and Madrid Senders. • Aggregates data. • Provides data to the TIBCO BusinessEvents Monitor displays.	• Automatically detects and gathers data from its local servers. • Sends data to the NYC Data Server.	• Automatically detects and gathers data from its local servers. • Sends data to the NYC Data Server.

Setting Up the Sender/Receiver Configuration

The following steps outline the workflow for setting up a Sender/Receiver configuration:

Receiver Configuration

This section assumes you have already installed the Monitor on the system where you will be running the receiver, and also that you have created a project directory. See Installation and Setup for information on installing the Monitor and Quick Start for an example of how to configure the Monitor.

Set up your receiver data server

1. Start the project using **start_server**. See Quick Start for more information.

2. By default, the receiver is setup to receive data on port 3272. If your senders cannot access the system on which the receiver is running, they can send data to the rtvagent servlet instead, which will forward the data to the receiver. To deploy the rtvagent servlet:

   If you are using Eclipse Jetty (the default application server):

   There are no required steps.

   If you are using Tomcat/a different application server:

   Copy the **rtview-tbemon-rtvagent.war** files located in the **TIB_rtview-be\projects\rtview-server** directory to the Tomcat **webapps** directory.

3. Remove any connections that will be serviced by a sender in the RTView Configuration Application > (**PROJECT NAME) > Solution Package Configuration** > **TIBCO BusinessEvents** > **CONNECTIONS**. If all connections will be serviced by senders, any connections created in the **CONNECTIONS** tab need to be removed. See Configuring Data Collection for more information.

4. Restart the project using **stop_server** and **start_server**. See Quick Start for more information.

Collector Configuration

This section assumes you have already installed the Monitor on the system where you will be running the sender, and also that you have created a project directory. See Installation and Setup for information on installing the Monitor and Quick Start for an example of how to configure the Monitor. You can run as many senders on as many systems as needed.

1. In the **rtvservers.dat** file located in your project directory, add **-propfilter:sender** to the end of the dataserver line and comment out the display server, historian and database processes as follows (since they are not used by sender data servers):

```
default . dataserver rundata -propfilter:sender
#default . historian runhist -ds
#default . displayserver rundisp -ds
#default . database rundb
```

2. Start the sender project using **start_server**. See Quick Start for more information.

   Note: If you are running multiple senders on the same system or running the sender on the same system as the receiver, you need to change the port prefix for the sender so that you do not get a port conflict. To do so, use the following on the command line as follows: **-portprefix:XX** where XX is the port prefix. To save this to your properties file so you do not need to specify it on the command line, add the **-saveportprefix** command line option. For example: **-portprefix:55 -saveportprefix**

3. Open the RTView Configuration Application > (**Project Name**) > **Server Configuration** > **Data Servers** > **COLLECTOR** tab.

4. In the **Targets** region, click the icon to add a target as follows:

   ID: A unique name for the target.

   URL: Specify the URL for the receiver. The url can be **host:port** (for example, somehost:3272) or an **http url** for the rtvagent servlet on the receiver. For example, if you are using Tomcat, you would use **http://somehost:8068/tbemon-rtvagent**. If you are using Jetty, you would use **http://somehost:3270/rtvagent**.

Targets: Select the **All solution packages** option.

Enabled: Select this check box to enable the target.

5. Click **Save** to exit the **Add Target** dialog.

6. Fill in a unique value for this sender in the **Identifier** field on the **COLLECTOR** tab. This should be unique across all senders.

7. Click on the RTView Configuration Application > (**Project Name**) > **Solution Package Configuration** > **TIBCO BusinessEvents** > **CONNECTIONS** tab and verify that this sender is configured to collect only from its local connections (see Configuring Data Collection for more information).

8. If you changed the port prefix in step 2, click on the RTView Configuration Application > (**Project Name**) > **Server Configuration** > **General** > **GENERAL** tab and confirm the port prefix is set to the correct value. If not, modify it accordingly.

9. Click **Save** in the RTView Configuration Application and restart your project using **stop_server** and **start_server**. See Quick Start for more information.

CHAPTER 5 Using the Monitor -- HTML Displays

This section describes how to read and use Monitor displays. This section includes:

- Overview
- TIBCO BusinessEvents Monitor HTML Views/Displays
- Drilldowns
- Alerts
- Admin

Overview

This section describes Monitor displays, how to read Monitor objects, GUI functionality and navigation. This section includes:

- TIBCO BusinessEvents Overview Display
- Tables
- Trend Graphs
- Time Settings
- Using the Monitor -- HTML Displays
- Export Report
- Rate Calculation in TIBCO BusinessEvents Monitor

TIBCO BusinessEvents Overview Display

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.

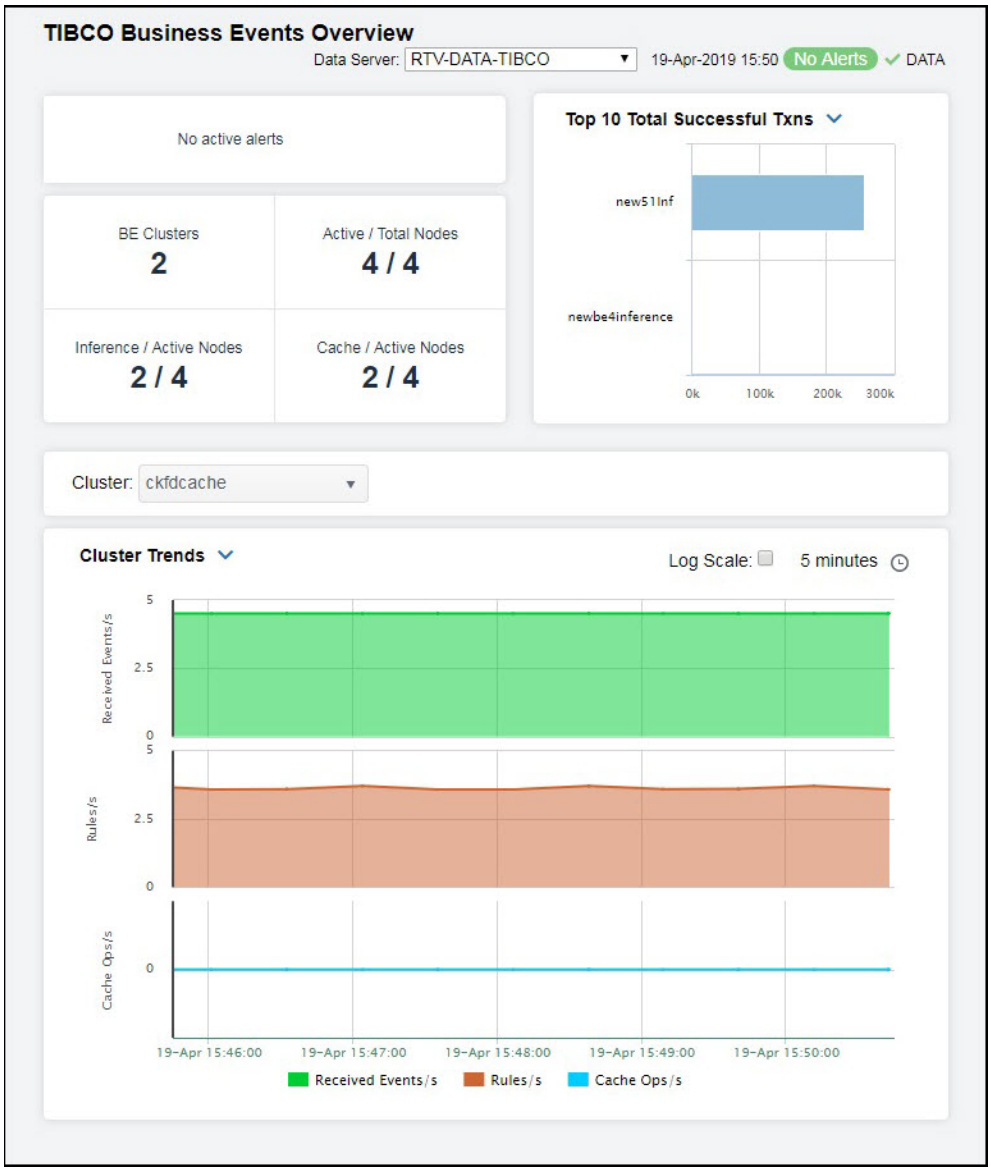

Tables

EMS Monitor tables contain the same data that is shown in the heatmap in the same View. Tables provide you a text and numeric view of the data shown in that heatmap, and additional

data not included the heatmap. For example, the **Cluster Nodes Table** display shows the same data as the **Cluster Nodes Heatmap** display.

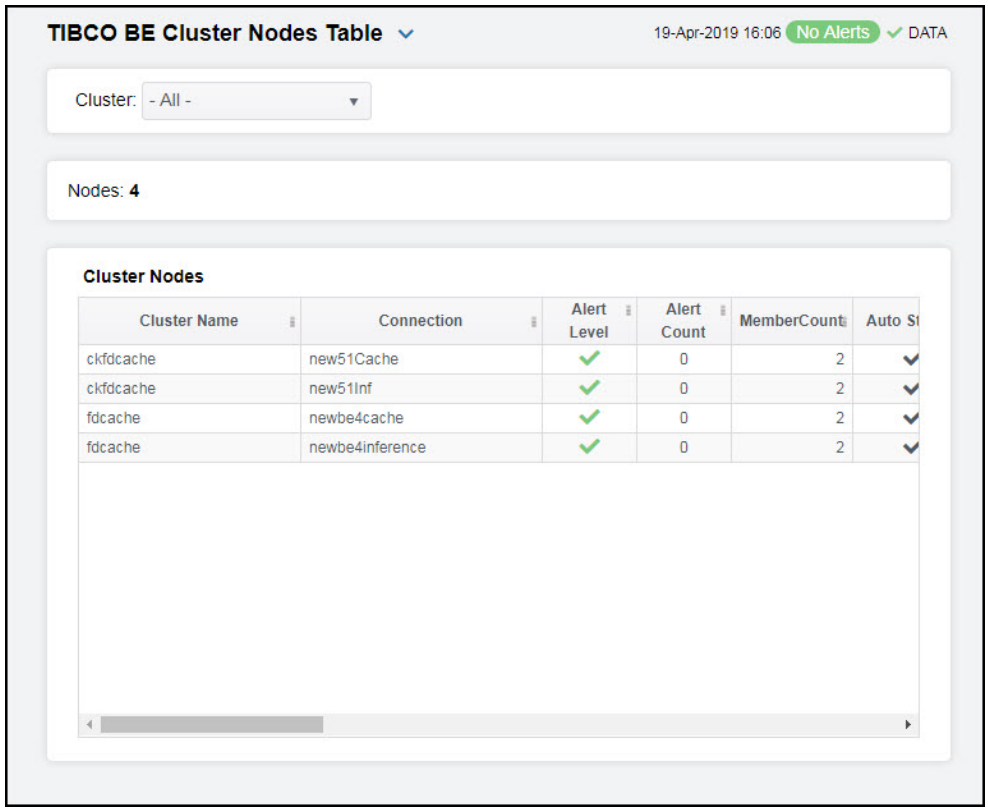

Table rows also sometimes use color to indicate the current most critical alert state for all resources associated with a given row. For example, the color coding is typically as follows:

● -- One or more alerts exceeded their critical threshold for one or more associated resources.

● --One or more alerts exceeded their warning threshold for one or more associated resources.

Tables support several interactive features: filtering on multiple columns, sorting, 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.

Additional features are:

- Column Filtering
- Multiple Column Sorting
- Column Visibility
- Column Reordering
- Row Paging

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:

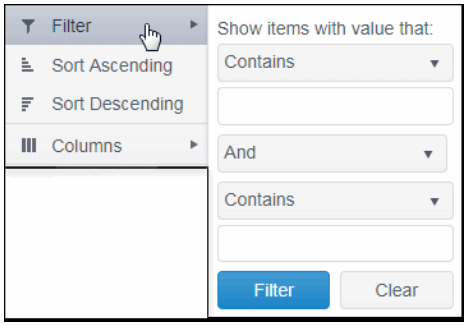

Options in the **Column Filter** dialog vary according to the data type of the selected column:

- **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.

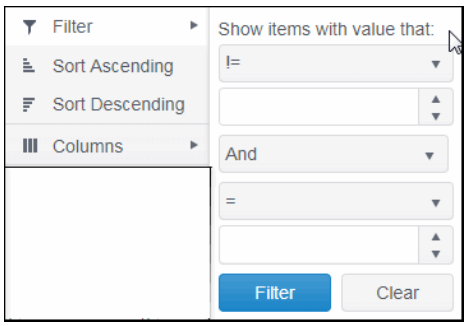

- **Date columns**: You can select a date and time and choose whether matching items should have a timestamp that is the same as, before, or after the filter time. The date is selected by clicking on the calendar icon and picking a date from a calendar dialog. The time is selected by clicking on the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

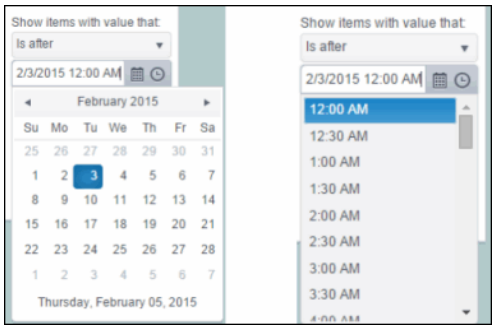

Alternatively, a date and time can be typed into the edit box. The strings shown in a date column are formatted by the Display Server using its time zone. But if a filter is specified on a date column, the date and time for the filter are computed using the client system's time zone. This can be confusing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

Multiple Column Sorting

Click on a column header and select **Sort Ascending** or **Sort Descending** to sort the table by that column. 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.

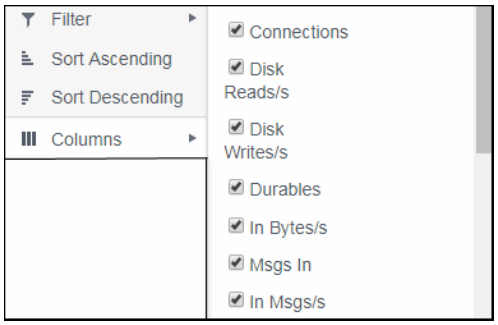

The leftmost column (the row header column) cannot be hidden.

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

Column Reordering

You can reorder the grid columns by dragging and dropping a column's header into another position.

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

Row Paging

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

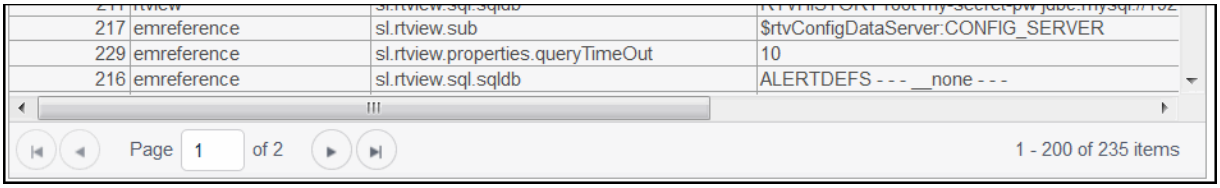

Trend Graphs

TIBCO BusinessEvents Monitor trend graphs enable you to view and compare performance metrics over time. You can use trend graphs to assess utilization and performance trends.

For example, the following figure illustrates a typical TIBCO BusinessEvents Monitor trend graph.

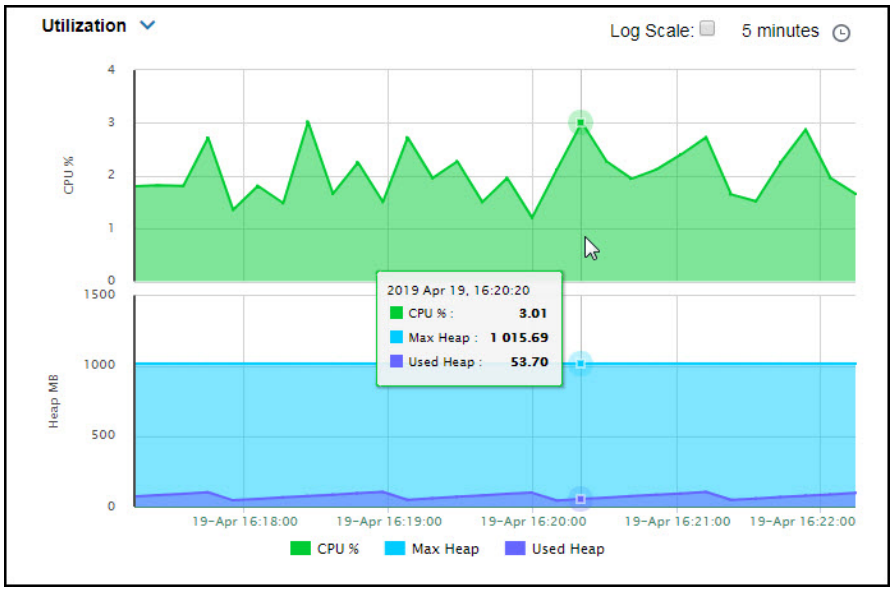

Time Settings

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.

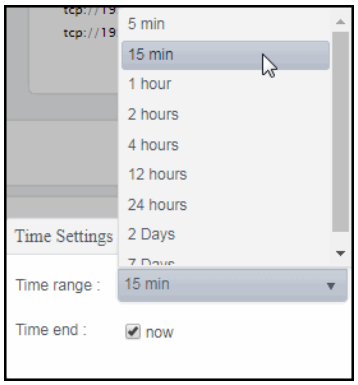

To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated 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.

Mouse-over

The mouse-over functionality provides additional detailed data in an over imposed pop-up window when you mouse-over trend graphs. The following figure illustrates mouse-over functionality. In this example, when you mouse-over a single dot, or data point, a pop-up window shows data for that data point.

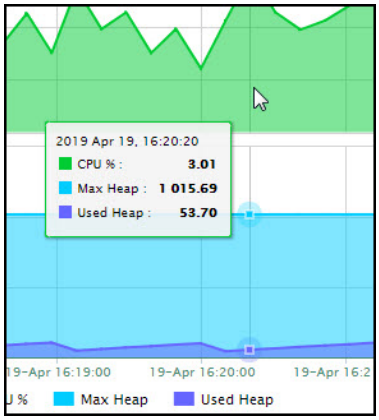

Log Scale

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.

Title Bar Functionality

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

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

	Displays the name of the data server for which you are viewing data.
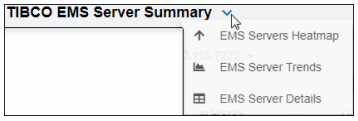	Navigates to a display that is most commonly accessed from the current display. The target display differs among displays.
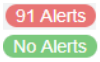	Displays the current number of alerts. Clicking on this "button" opens the Alerts Table display in a new window.
	The current date and time. If the time is incorrect, this might indicate that RTView stopped running. When the date and time is correct and the **Data OK** indicator is green, this is a strong indication that the platform is receiving current and valid data.
	The data connection state. Red indicates the data source is disconnected (for example, if the Data Server is not receiving data, or if the Display Server does not receive data from the Data Server, this will be red). Green indicates the data source is connected. When the date and time is correct and the **Data OK** indicator is green, this is a strong indication that the platform is receiving current and valid data.

Export Report

You can quickly export reports for tables in a display by right-clicking on the display and selecting **Export to Excel > Current Page or All Pages**.

Rate Calculation in TIBCO BusinessEvents Monitor

The majority of rate calculations in BE Monitor use the following property, which is defined in the **\rtvapm\tbemon\conf\rtvapm.tbemon.properties** file:

```
#Query Interval

#

sl.rtview.sub=$tbeQueryIntervalBase:30
```

For **Events Received Per Sec**, for example, if the query interval was defined as 10 (seconds), let's say you had the following:

Query (Time) Interval	Total Number of Events Received
t0	10
t0+10 (seconds)	20
t0+20 (seconds)	30

In the table above, you can see that 10 total events were received during the first query interval. After the second query interval, 20 total events were received. After the third query interval, 30 total events were received. The rate of events received per second for the current query interval is then calculated by taking the total number of events received after the current query interval, subtracting the total number of events received after the previous query interval (to determine the delta), and dividing by the query interval. For example:

Query (Time) Interval	Total Number of Events Received after Query Interval	Delta Time	Delta Number of Events Received (current - previous)	Events Received Per Sec
t0	10	-	-	-
t0+10 (seconds)	20	10	(20-10)	(20-10)/10= 1 Event/sec
t0+20 (seconds)	30	10	(30-20)	(30-20)/10= 1 Event/sec

So, **Events Received Per Sec** is calculated by taking the difference in the number of Events (delta) between the current and previous query interval (30 events - 20 events) and dividing by the current query interval (10 seconds), or:

(30-20)/10 (seconds)= 10/10 = 1 Event per second

TIBCO BusinessEvents Monitor HTML Views/Displays

The TIBCO BusinessEvents Monitor contains the following Views:

- BE Clusters - HTML
- BE Nodes - HTML
- BE Events - HTML
- BE Concepts - HTML

BE Clusters - HTML

These displays present performance metrics and alert status for your BusinessEvents system. Clicking **BE Clusters** from the left/navigation menu opens the TIBCO BE Clusters Table - HTML display, which shows all available utilization metrics for all BE clusters. The options available under **BE Clusters** are:

- **Clusters Heatmap**: Opens the TIBCO BE Clusters Heatmap - HTML, which shows cluster and alert status for all BE clusters.
- **Clusters Summary**: Opens the TIBCO BE Cluster Summary - HTML display, which shows information for a single BE Cluster.

TIBCO BE Clusters Table - HTML

Use this display to check event, concept, and backing store metrics for all of your clusters. Consider keeping this display open to monitor your TIBCO BusinessEvents clusters in general. Each row in the table contains data for a particular cluster. Click a column header to sort column data in ascending or descending order. Double-click on a table row to drill-down to the **TIBCO BE Cluster Summary - HTML display and view metrics for that particular cluster. Toggle between the commonly accessed Table and Heatmap displays by clicking the drop down list on the display title.**

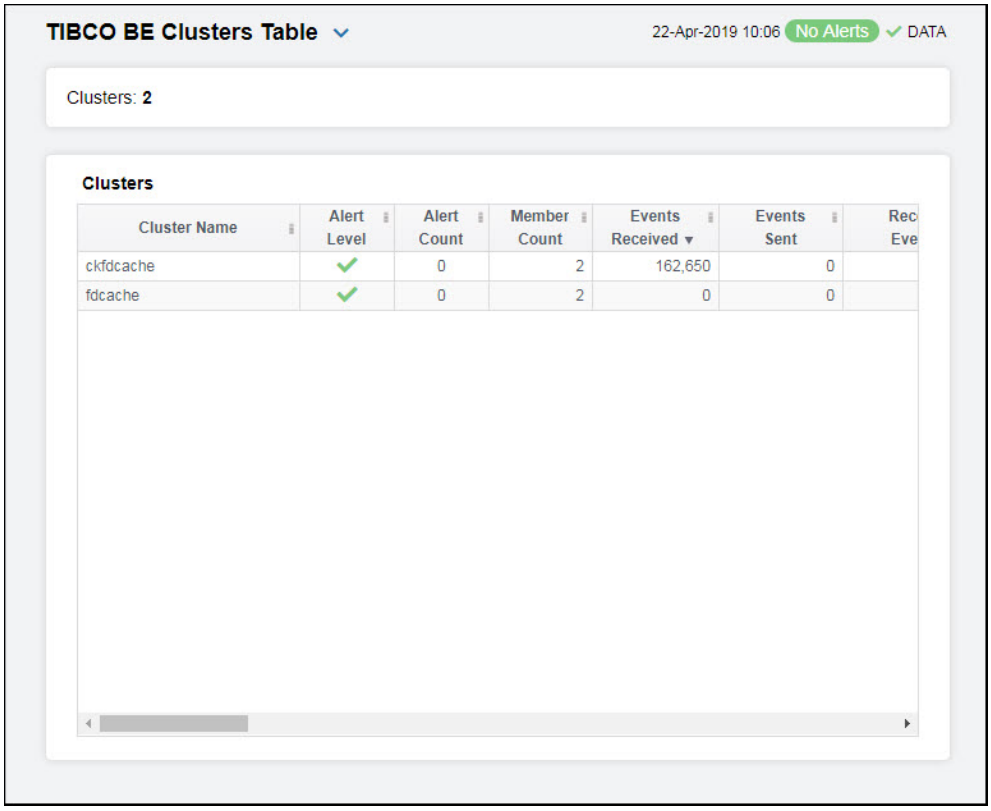

Clusters Table

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	The rate of "loads" into the Backing Store.

Loads/s	
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 - HTML

Clicking **Clusters Heatmap** in the left/navigation menu opens the **TIBCO BE Clusters Heatmap**, which allows you to view the status and alerts of all BE clusters. Use the **Metric** drop-down menu to view the heatmap using a different metric.

The heatmap is organized so that each rectangle represents a cluster. The rectangle color indicates the most critical alert state. Click on a node to drill-down to the TIBCO BE Cluster Summary - HTML display and view metrics for a particular cluster. Toggle between the commonly accessed **Table** and **Heatmap** displays by clicking the drop down list on the display title. Mouse-over rectangles to view more details about cluster performance and status.

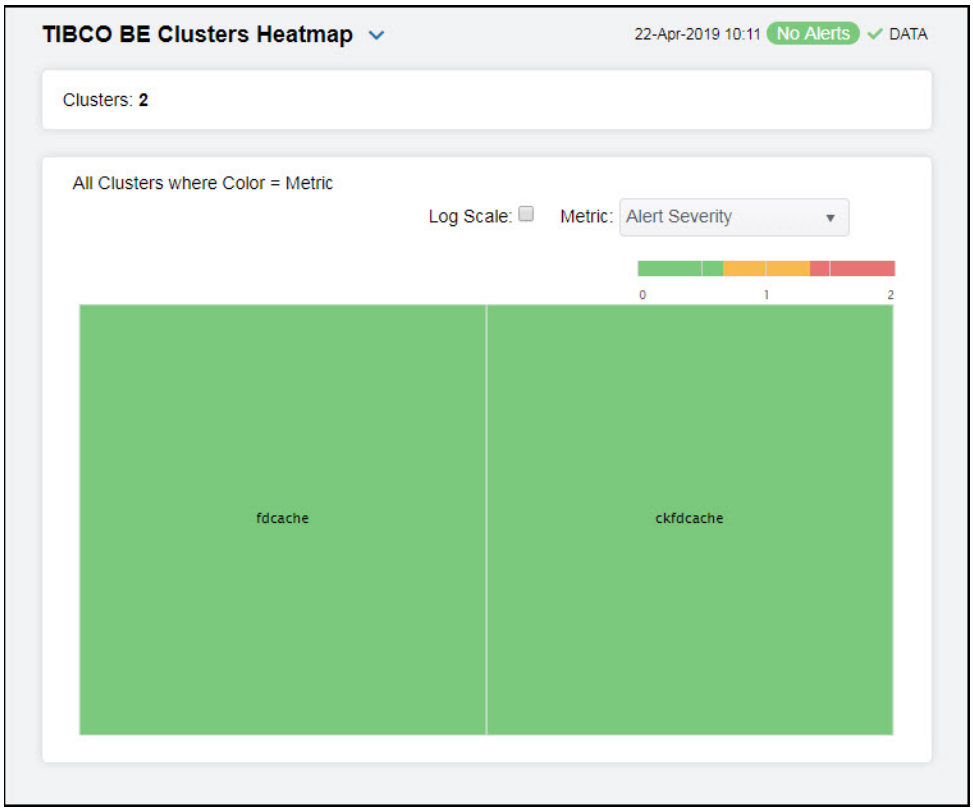

Fields and Data

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 - HTML display for a detailed view of metrics for that particular cluster.

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

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

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

Member Count The total number of members in the cluster. The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum number of members in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

Events Received The number of events received in the cluster. The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the most events received in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

Events Sent The number of events sent in the cluster. The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the most events sent in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

TIBCO BE Cluster Summary - HTML

Clicking **Cluster Summary** in the left/navigation menu opens the **TIBCO BE Cluster Summary** display, which allows you to view configuration and utilization data for a single cluster. Select a cluster to view Rete statistics, cache metrics, Backing Store data, and trend data for the cluster. Clicking on the information boxes at the top of the display takes you to the TIBCO BE Clusters Table - HTML display or the TIBCO BE Events Table - HTML display, where you can

view additional cluster/event data. There are two options in the trend graph region: **Cluster Trends** and **Backing Store Trends**. In the **Cluster Trends** trend graph region, you can view received event rate, rules fired rate, and cache operations rate over a selected time range. In the **Backing Store Trends** trend graph region, you can view backing store operations rate and average maximum time per backstore operation over a selected time range. Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display.

F**ilter 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	**Cluster Trends** Shows the following metrics for the selected cluster. **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. **Backing Store Trends** Shows the following metrics for the selected cluster. **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.
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. 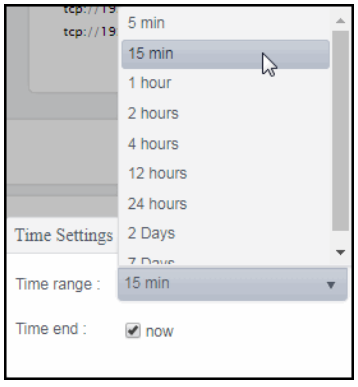 To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated 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.

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 - HTML

These displays present performance metrics and alert status for your BusinessEvents nodes. Clicking **BE Nodes** from the left/navigation menu opens the TIBCO BE Cluster Nodes Table - HTML display, which shows all available utilization metrics for all BE nodes. The options available under **BE Nodes** are:

- **Cluster Nodes Heatmap**: Opens the TIBCO BE Cluster Nodes Heatmap - HTML, which shows cluster and alert status for all BE cluster nodes.
- **Inference Node Summary**: Opens the TIBCO BE Inference Node Summary - HTML display, which shows information for (inference) agents for a single BE cluster node.
- **Storage Node Summary**: Opens the TIBCO BE Storage Node Summary - HTML display, which displays cache data for a specific node.

TIBCO BE Cluster Nodes Table - HTML

Use this display to view configuration and utilization data for nodes in a cluster. Each row in the table contains data for a particular node. Click a column header to sort column data in ascending or descending order. Toggle between the commonly accessed **Table**, **Heatmap**, and **Summary** displays by clicking the drop down list on the display title.

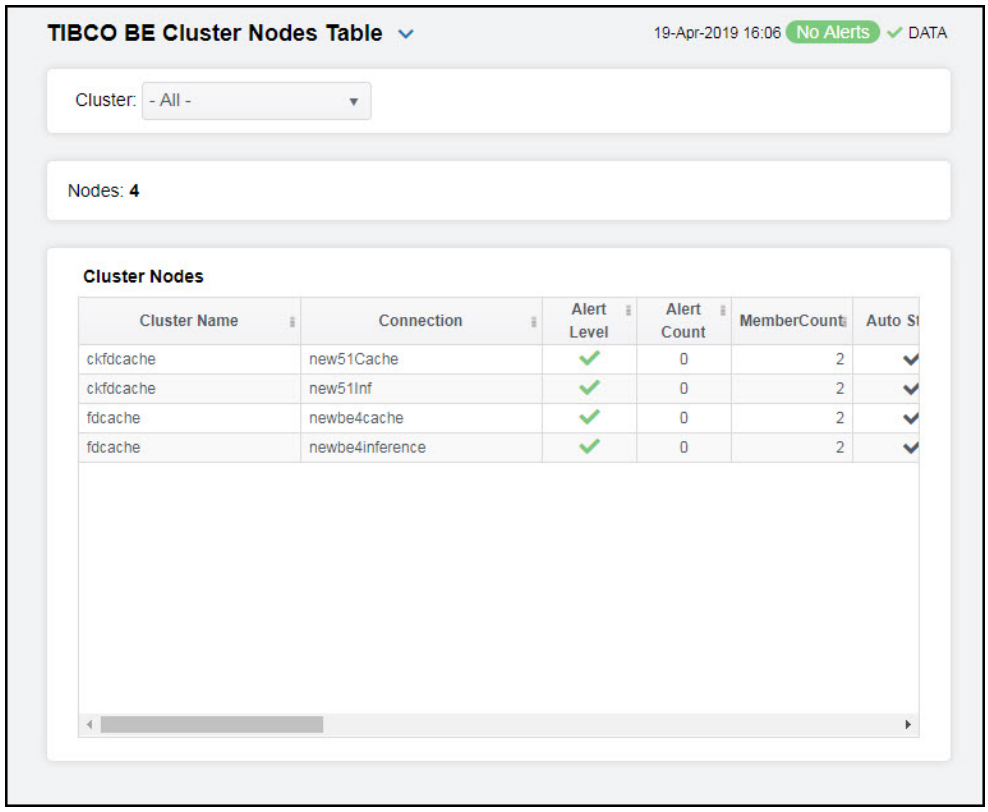

Filter By:

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 - HTML

Clicking **Cluster Nodes Heatmap** in the left/navigation menu opens the **TIBCO BE Cluster Nodes Heatmap**, which allows you to view utilization data for all nodes in a cluster in a heatmap format. Use the **Metric** drop-down menu to view the heatmap based on a different metric.

The heatmap is organized so that each rectangle represents a node. The rectangle color indicates the most critical alert state. Toggle between the commonly accessed **Table** and **Heatmap** displays by clicking the drop down list on the display title. Mouse-over rectangles to view more details about cluster performance and status.

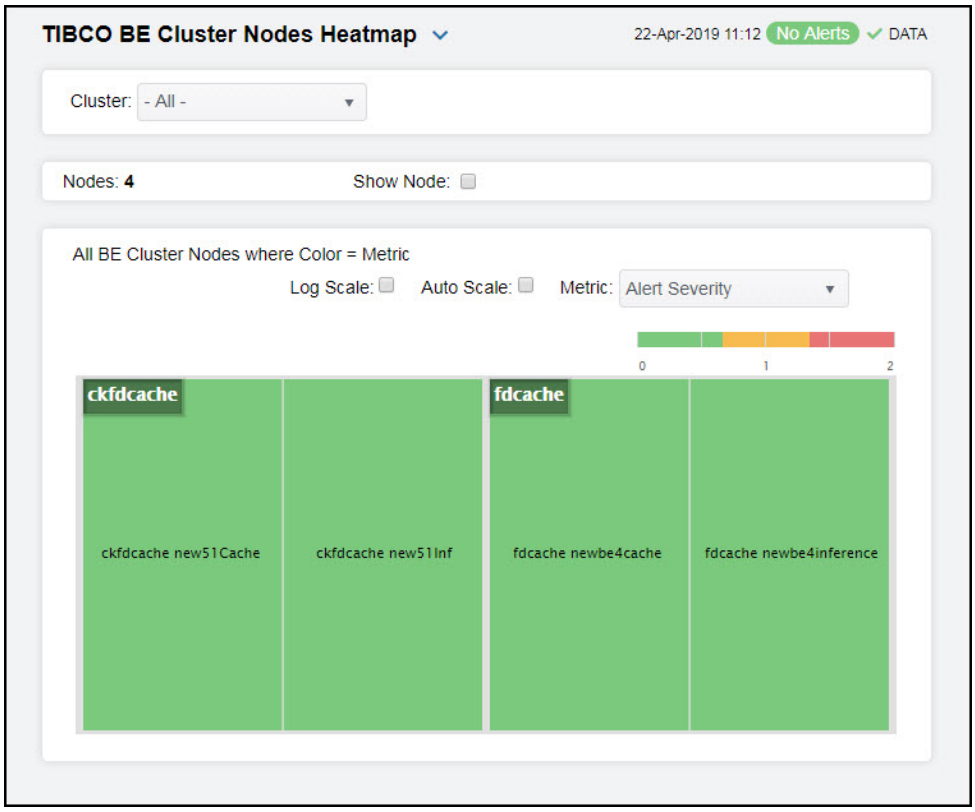

Filter By:

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 The maximum alert level in the item (index) associated with the rectangle. Values range from **0** to **2**, as indicated in the color gradient bar 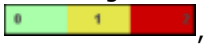 , where **2** is the greatest **Alert**

	Severity. **2** -- Metrics that have exceeded their specified **ALARMLEVEL** threshold and have an Alert Severity value of **2** are shown in red. For a given rectangle, this indicates that one or more metrics have exceeded their alarm threshold. **1** -- Metrics that have exceeded their specified **WARNINGLEVEL** threshold and have an Alert Severity value of **1** are shown in yellow. For a given rectangle, this indicates that one or more metrics have exceeded their warning threshold. **0** -- Metrics that have not exceeded either specified threshold have an Alert Severity value of **0** and are shown in green. For a given rectangle, this indicates that no metrics have exceeded a specified alert threshold.
Alert Count	The total number of alarm and warning alerts in a given item (index) associated with the rectangle. The color gradient bar 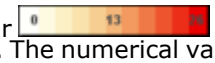 shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the middle value of the range.
JVM % CPU Used	The total percentage of JVM CPU used in a given item (index) associated with the rectangle. The color gradient bar 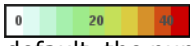 shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of **JvmCpuPercentHigh**. The middle value in the gradient bar indicates the middle value of the range.
JVM % Heap Used	The total percentage of JVM Heap Memory Used in a given item (index) associated with the rectangle. The color gradient bar 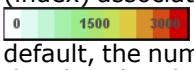 shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of **JvmMemoryUsedHigh**. The middle value in the gradient bar indicates the middle value of the range.

TIBCO BE Inference Node Summary - HTML

Clicking **Inference Node Summary** in the left/navigation menu opens the **TIBCO BE Inference Node Summary** display, which allows you to view configuration and utilization data for a single inference node. View a list of all agents on the node and trend graphs tracing the rule execution rate for agents on the node. The rule execution rate is relative to the overall CPU and heap utilization for the engine's JVM.

NOTE: An inference node (also known as an engine or processing unit) is the container where one or more inference agents run. Generally, the agents in a given node implement different rule sets, and distributing nodes on different hosts provides fault tolerance and load balancing for the cluster. For details, refer to TIBCO documentation.

Clicking on the information boxes at the top of the display takes you to the TIBCO BE Clusters Table - HTML display, the TIBCO BE Agent Event Summary - HTML display, or the **JVM Summary** display, where you can view additional cluster/event/JVM data. The **Agents for this Node** region lists the inference agents associated with the selected node. There are two options in the trend graph region: **Utilization** and **Rules/s and Threads**. In the **Utilization** trend graph region, you can view percentage of CPU being used by the engine/process and the amount of memory, in megabytes, in use by the JVM for heap space over a selected time range. In the **Rules/s and Threads** trend graph region, you can view the rate of rules fired for the agent and the total number of threads for the agent over a selected time range. Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display.

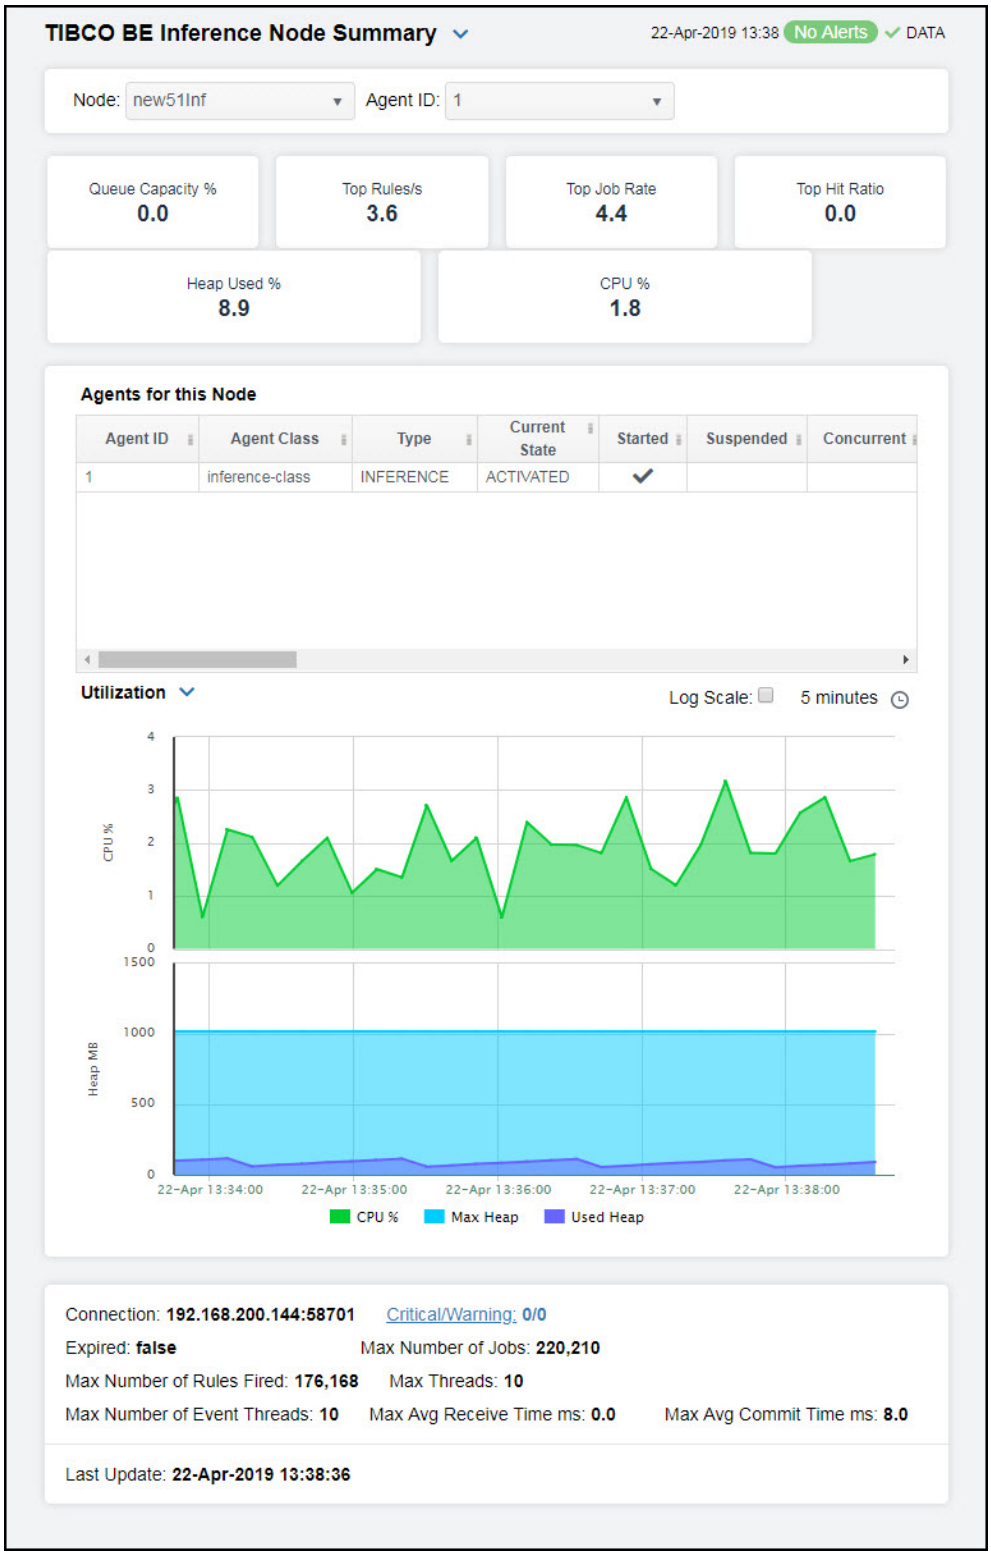

Filter By:
The display might include these filtering options:

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. 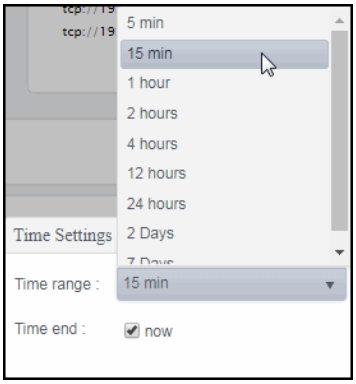 To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the

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 - HTML

Clicking **Storage Node Summary** in the left/navigation menu opens the **TIBCO BE Storage Node Summary** display, which allows you to view configuration details for a single cache node and a list of all caches that are backed by the backing store (database).

NOTE: A storage node (also known as a cache node) provides fast access to events and concepts required during each RTC by the inference engines. Storage nodes also serve as buffers for reads and writes between the cluster and the backing store. For details, refer to TIBCO documentation.

Clicking on the information boxes at the top of the display takes you to the **JVM Summary** display, where you can view additional JVM data. The **Backing Store** region lists the caches that are backed by the backing store in the selected node. There are two options in the trend graph region: **Utilization** and **Table Sizes**. In the **Utilization** trend graph region, you can

view percentage of CPU being used by the node and the amount of memory, in megabytes, in use by the JVM for heap space over a selected time range. In the **Table Sizes** trend graph region, you can view the number of unique objects cached in the local index table and the number of entries in the table of external IDs used as indexes by the backing store over a selected time range. Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display.

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.

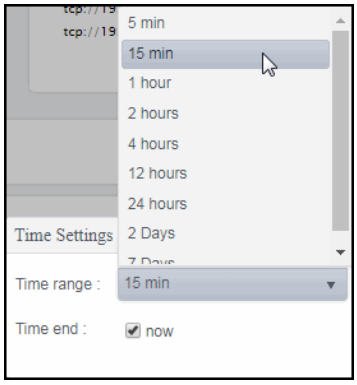

To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated 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 - HTML

These displays present performance metrics and alert status for your BusinessEvents events. Clicking **BE Events** from the left/navigation menu opens the TIBCO BE Events Table - HTML display, which shows all available utilization metrics for all BE events. The options available under **BE Events** are:

- **Event Summary**: Opens the TIBCO BE Agent Event Summary - HTML display, which shows information for a single BE event.
- **Event Cache Hits**: Opens the TIBCO BE Event Cache Hits Table - HTML display, which
- **Event Hit Summary**: Opens the TIBCO BE Event Hit Summary - HTML display, which

TIBCO BE Events Table - HTML

View run-time statistics for a selected group of agents. With TIBCO BusinessEvents, events are cached when they are out of scope, and deleted or persisted to the backing store when they are no longer useful. Double-clicking on a row in the table displays access patterns over time for the event in the TIBCO BE Agent Event Summary - HTML display. Each row in the table contains data for a particular event. Click a column header to sort column data in ascending or descending order. Toggle between the commonly accessed **Table** and **Summary** displays by clicking the drop down list on the display title.

NOTE: Events cause rules to execute in the BusinessEvents Rete network. Events can be created by external phenomena, such as the arrival of a JMS message, or internally when rules are processed. When an event enters the Rete network, it causes a run-to-completion cycle which continues until no further rules can be processed. Each named event that can be handled by a BusinessEvents application is specified at build time in BusinessEvents studio. For details, refer to TIBCO documentation.

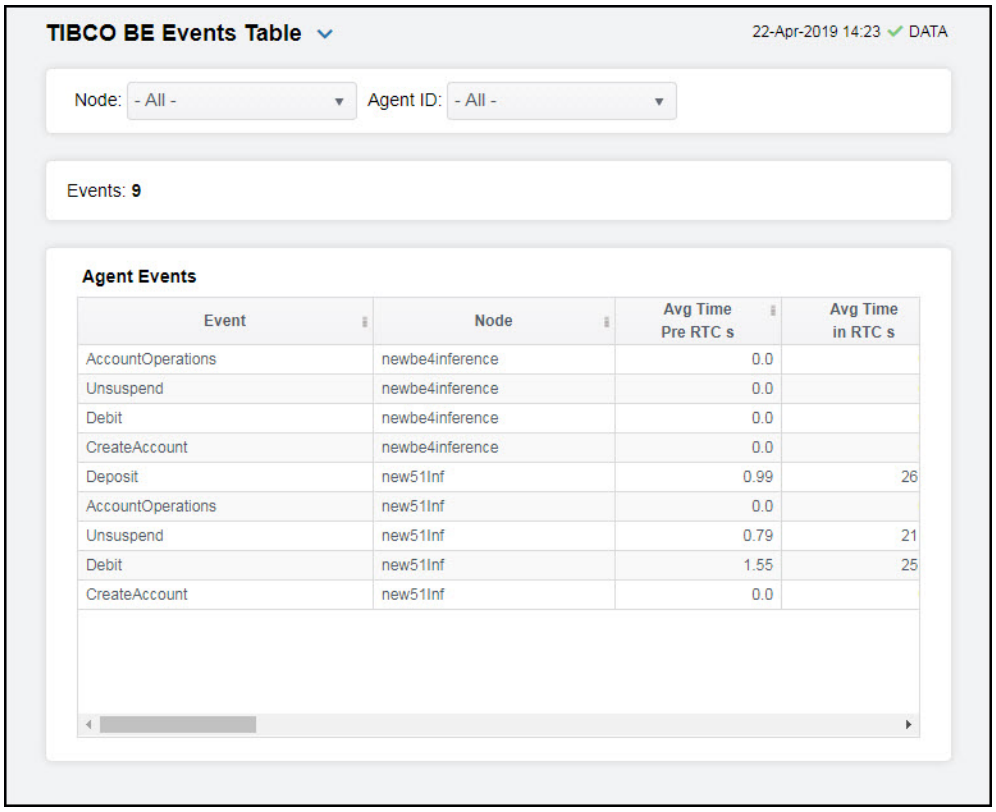

Filter By:

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	The number of times data has been searched for in the L1 cache, but was not

Cache	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 - HTML

Clicking **Event Summary** in the left/navigation menu opens the **TIBCO BE Agent Event Summary** display, which allows you to view detailed performance metrics for an agent's event.

Clicking on the information boxes at the top of the display takes you to the TIBCO BE Events Table - HTML display, where you can view additional event data. There are three options in the trend graph region: **Time in RTC**, **Agent Rates**, and **Channel Rates**. In the **Time in RTC** trend graph region, you can view the average amount of time taken for the event to begin its run to completion cycle, the average amount of time taken for the event to complete (once it has started) its run to completion cycle, and the average amount of time taken by the event after its run to completion cycle has ended over a selected time range. In the **Agent Rates**

trend graph region, you can view rate of event assertions into the Rete network via the agent, the rate of event retractions/deletions from the Rete network via the agent, and the rate of event modifications in the Rete network via the agent over a selected time range. In the **Channel Rates** trend graph region, you can view rate of event assertions into the Rete network via the channel, the rate of event retractions/deletions from the Rete network via the channel, and the rate of event modifications in the Rete network via the channel over a selected time range. Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display.

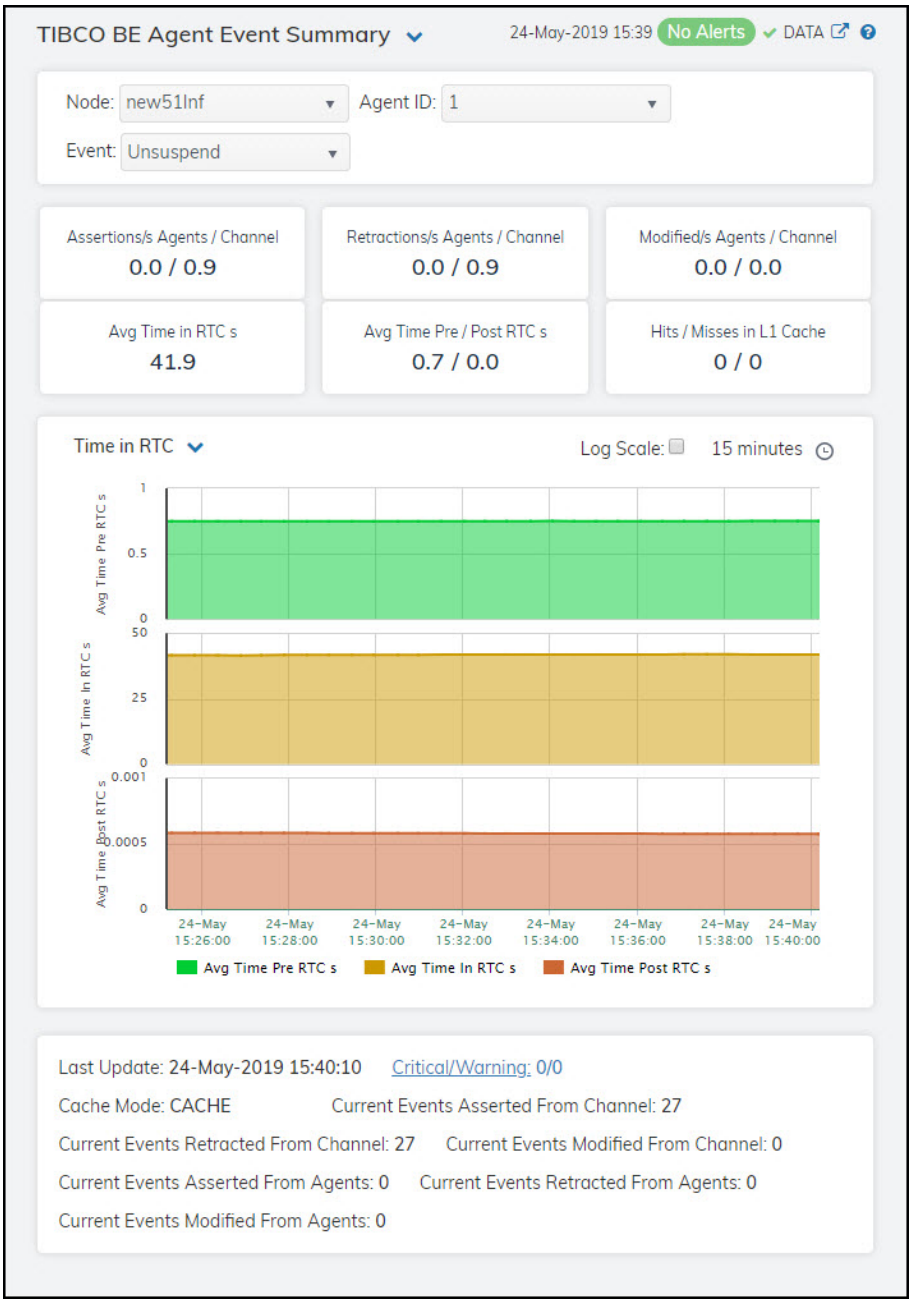

Filter By:

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.

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.

Time Settings

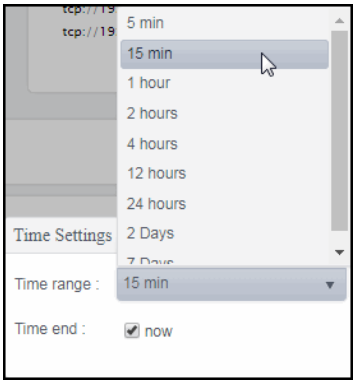

To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated 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	The number of event modifications in the Rete network via the agents.

Modified From Agents

TIBCO BE Event Cache Hits Table - HTML

Clicking **Event Cache Hits** in the left/navigation menu opens the **TIBCO BE Event Cache Hits Table** display, which allows you to view cache performance metrics per event for a single cluster or **All Clusters**.

NOTE: Events cause rules to execute in the BusinessEvents Rete network. Events can be created by external phenomena, such as the arrival of a JMS message, or internally when rules are processed. When an event enters the Rete network, it causes a run-to-completion cycle which continues until no further rules can be processed. Each named event that can be handled by a BusinessEvents application is specified at build time in BusinessEvents studio. For details, refer to TIBCO documentation.

Click a column header to sort column data in ascending or descending order. Double-click on a table row to drill-down to the **TIBCO BE Event Hit Summary - HTML** **display and view metrics for that particular event. Toggle between the commonly accessed Table and Summary displays by clicking the drop down list on the display title.**

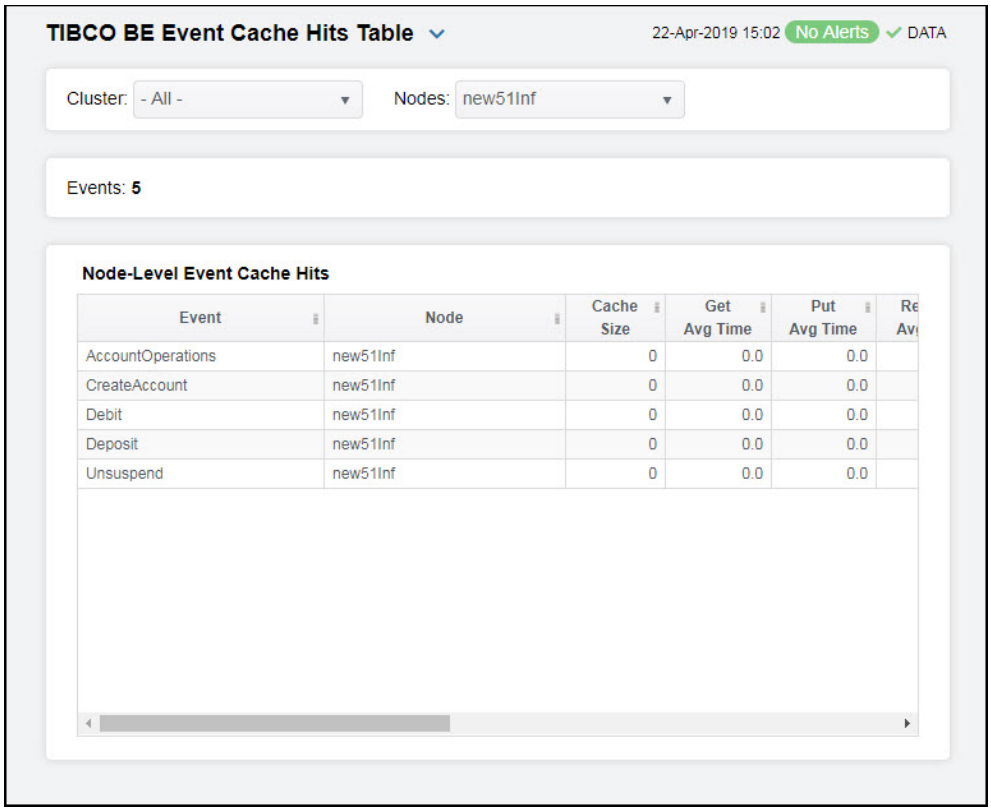

Filter By:

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 - HTML

Clicking **Event Hit Summary** in the left/navigation menu opens the **TIBCO BE Event Hit Summary** display, which allows you to view detailed event performance metrics for a single cluster or **All Clusters,** a node, and an event.

NOTE: Events cause rules to execute in the BusinessEvents Rete network. Events can be created by external phenomena, such as the arrival of a JMS message, or internally when rules are processed. When an event enters the Rete network, it causes a run-to-completion cycle which continues until no further rules can be processed. Each named event that can be handled by a BusinessEvents application is specified at build time in BusinessEvents studio. For details, refer to TIBCO documentation.

There are two options in the trend graph region: **Hit Rates** and **Current Hits**. In the **Hit Rates** trend graph region, you can view the number of "get" operations per second, the number of "put" operations per second, and the number of "remove" operations per second over a selected time range. In the **Current Hits** trend graph region, you can view the number of "get" operations, the number of "put" operations, and the number of "remove" operations

over a selected time range. Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display.

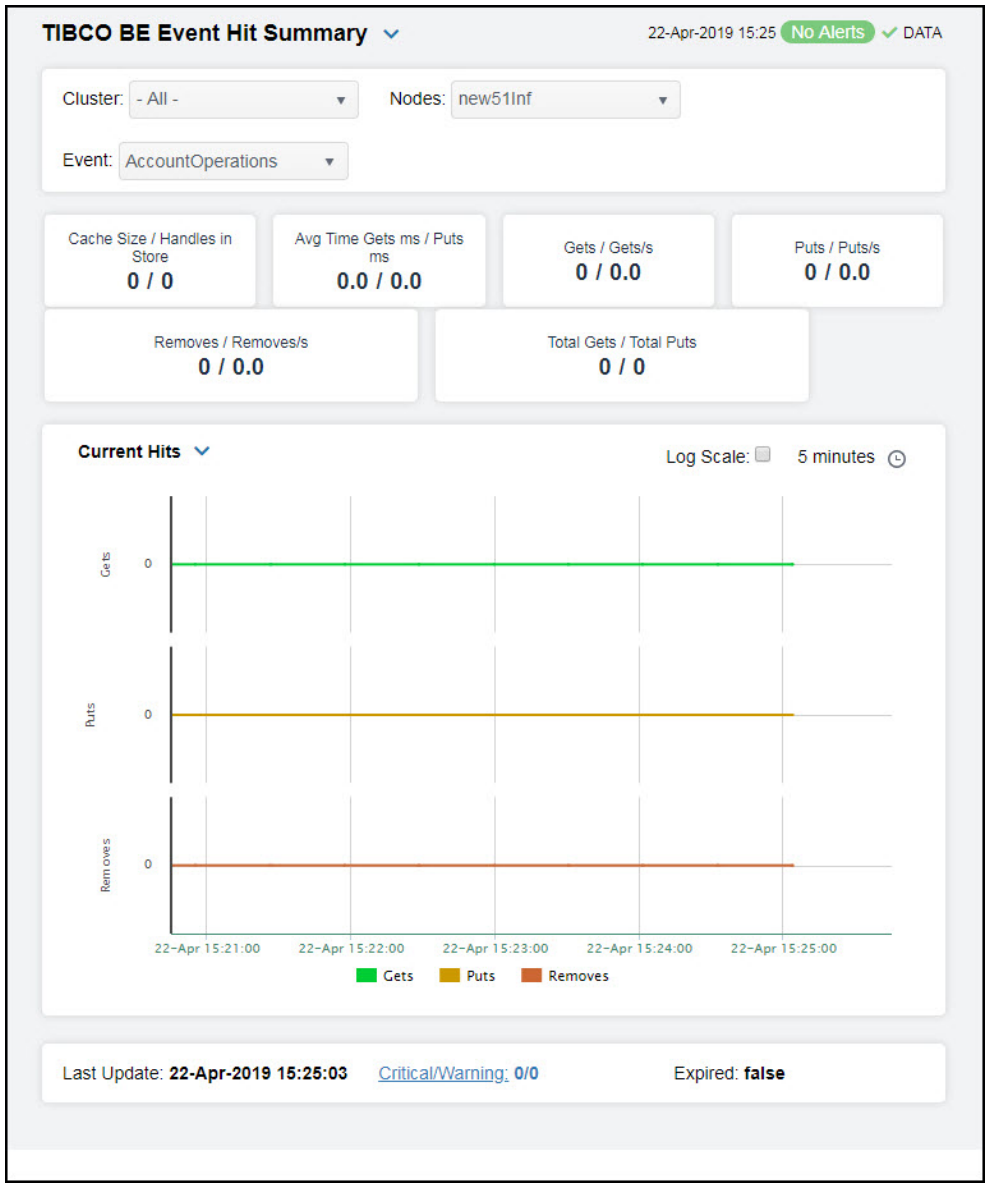

Filter By:

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	The size of the cache, and the number of handles in the Backing Store

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.

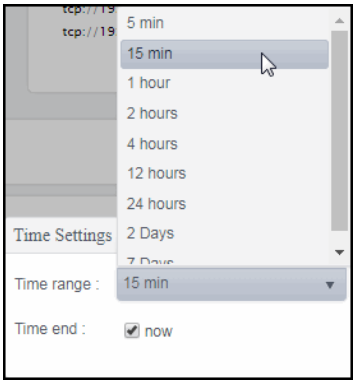

To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per click, or you can choose the date and time from the associated 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 - HTML

These displays present performance metrics and alert status for your BusinessEvents concepts. Clicking **BE Concepts** from the left/navigation menu opens the TIBCO BE Concepts Table - HTML display, which shows all available utilization metrics for all BE concepts. The options available under **BE Concepts** are:

- **Concept Hit Summary**: Opens the TIBCO BE Concept Hit Summary - HTML, which shows details and alert status for a BE concept.
- **Channels**: Opens the TIBCO BE Channel Status Table - HTML display, which shows information for destinations and channels for a single BE node.
- **Inference Agents Table**: Opens the TIBCO BE Inference Agents Table - HTML display, which displays agents data for a specific cluster.
- **RTC Reports Table**: Opens the TIBCO BE RTC Txn Manager Reports - HTML display, which displays reports data for a specific cluster.

TIBCO BE Concepts Table - HTML

View a list of concepts and their run-time statistics. Choose a single cluster or **All Clusters** and a node from the drop-down menus. Double-clicking on a row in the table displays additional details as well as hit rates and current hits over time for the concept in the TIBCO BE Concept Hit Summary - HTML display. Each row in the table contains data for a particular concept. Click a column header to sort column data in ascending or descending order. Toggle between the commonly accessed **Table** and **Summary** displays by clicking the drop down list on the display title.

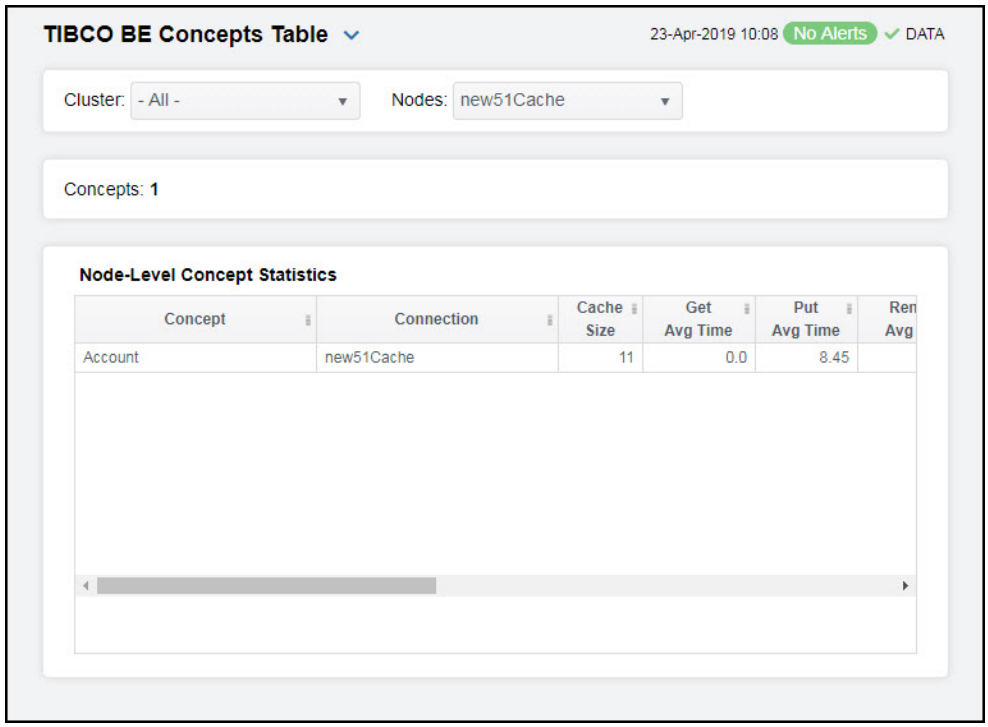

Filter By:

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 - HTML

Clicking **Concept Hit Summary** in the left/navigation menu opens the **TIBCO BE Concept Hit Summary** display, which allows you to view current and historic data for a single concept. Data in this display can be useful if your BusinessEvents system uses Cache object management. When Cache object management is used, concepts with a sufficiently long time to live (TTL) setting are cached.

Cache reference patterns for certain concepts may be related to incoming events (for example, customer purchase orders with associated inventory queries). The trend charts show the cache activity of such concepts, and might be useful in diagnosing the behavior of your application over time.

There are two options in the trend graph region: **Hit Rates** and **Current Hits**. In the **Hit Rates** trend graph region, you can view the number of "get" operations per second, the number of "put" operations per second, and the number of "remove" operations per second over a selected time range. In the **Current Hits** trend graph region, you can view the number of "get" operations, the number of "put" operations, and the number of "remove" operations over a selected time range. Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display.

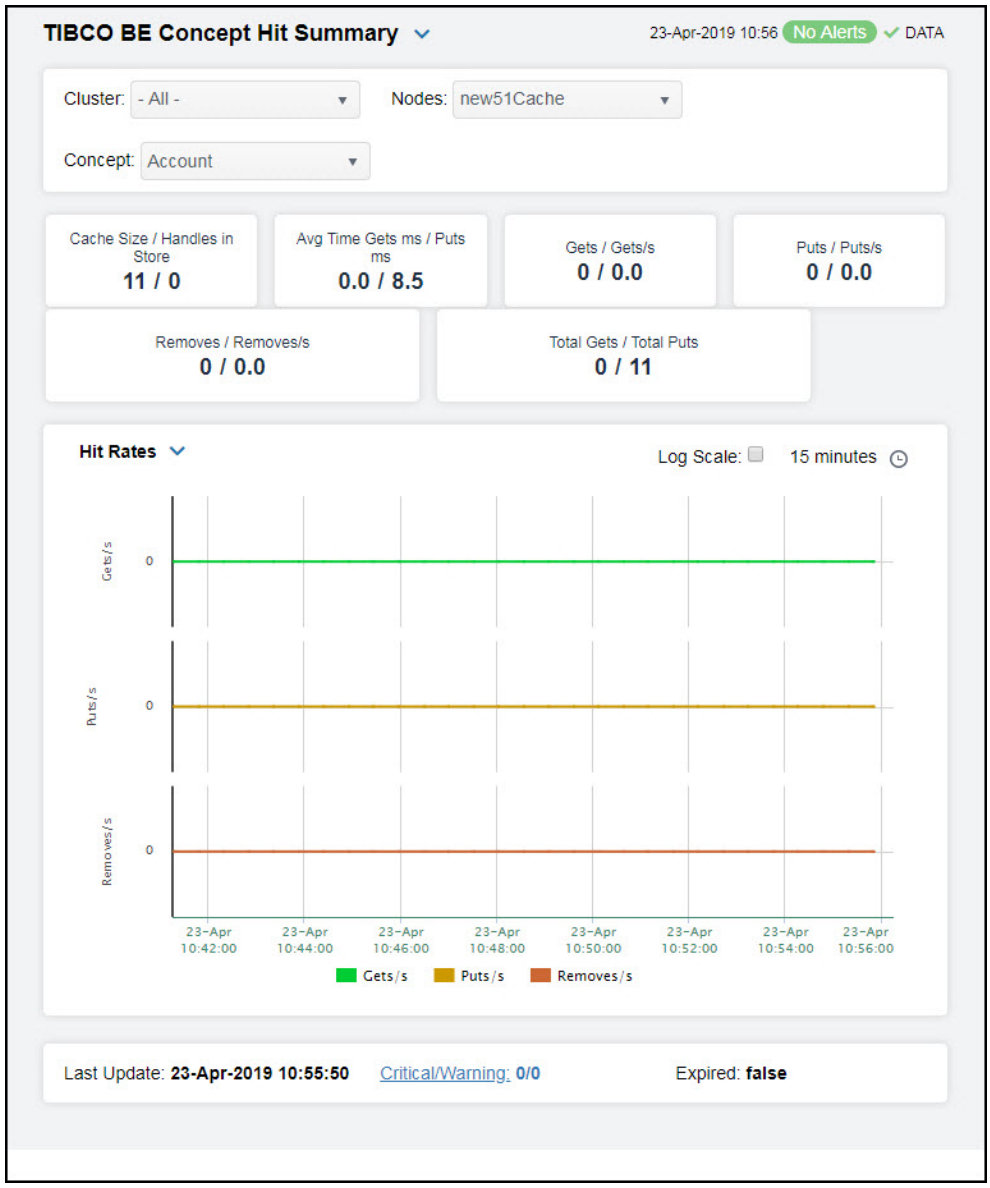

Filter By:

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.

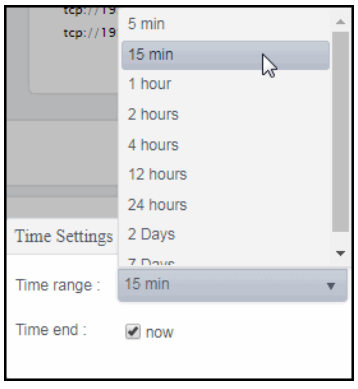

To change the time range, deselect the **now** toggle, which displays some additional date fields. You can click the left and right arrow buttons to decrease the end time by one time period (the time selected in the **Time range** drop down) per

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 - HTML

Clicking **Channels** in the left/navigation menu opens the **TIBCO BE Channel Status Table** display, which allows you to view a list of destinations, which are sources and sinks of events, and Channels. Destinations are potentially bi-directional, and the table indicates whether events are sent or received. Channels provide a class wrapper for destinations, and make it possible to enable or disable a group of destinations with one operation.

Click a column header to sort column data in ascending or descending order. Toggle between the commonly accessed **Table** and **Summary** displays by clicking the drop down list on the display title.

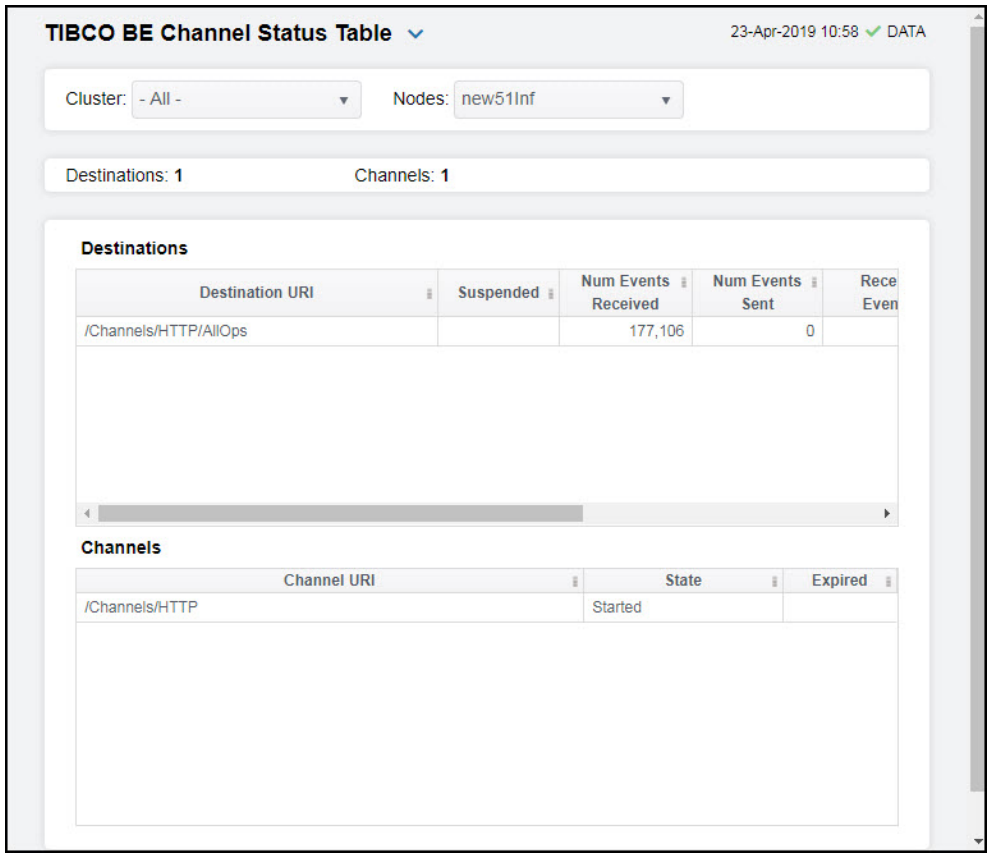

Filter By:

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 able 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 able 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 - HTML

Clicking **Inference Agents Table** in the left/navigation menu opens the **TIBCO BE Inference Agents Table** display, which allows you to compare agent metrics across deployed engines and verify that the cluster is properly load-balanced. View a list of all the inference agents deployed in each cluster. You can view agent data for a single cluster or all clusters.

The data in this display is identical to the data provided for a single engine in the TIBCO BE Cluster Summary - HTML display, except that it is aggregated across all inference nodes.

Choose a single cluster or **All** clusters from the drop-down menu. Each row in the table is a different agent.

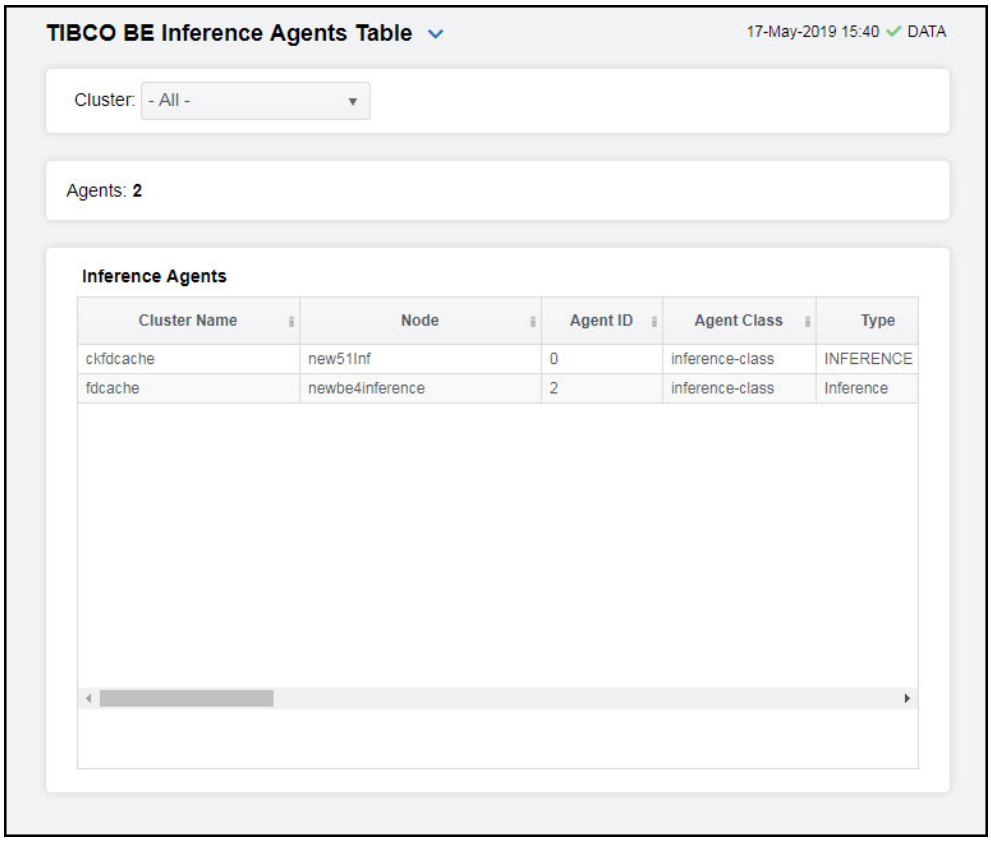

Filter By:

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 - HTML

Clicking **RTC Reports Table** in the left/navigation menu opens the **TIBCO BE RTC Txn Manager Reports** display, which allows you to compare RTC metrics across deployed engines. View a list of all the inference engine RTC reports. You can view reports for a single cluster or all clusters. The data in this display is identical to the data provided for a single engine in the TIBCO BE Cluster Summary - HTML display, except that it is aggregated across all inference nodes.

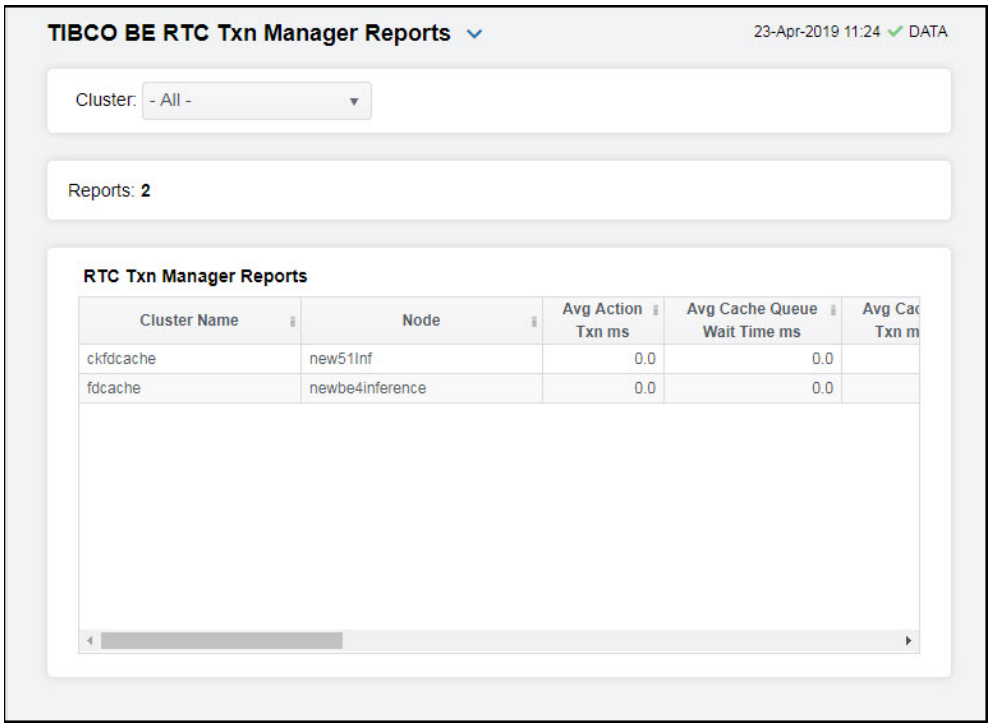

Filter By:

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	The size of the last database batch.*

Batch Size	
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.

Drilldowns

Displays contained under this View are only available by clicking on buttons/links in other displays. For example, clicking on the "Alerts" icon

in the upper right hand corner of the display opens the **Alerts Table by Component** display. You cannot access these displays directly via the left menu.

- **Component Alerts Table**: Associated with the Alerts Table by Component display, which can be accessed by clicking the Alerts icon

in the upper right hand corner of the display.
- **Component Alert Detail**: Associated with the Alert Detail for Component display, which can be accessed by clicking

in the **Alerts Table** or by clicking

in the **Alerts Table by Component** display.

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.

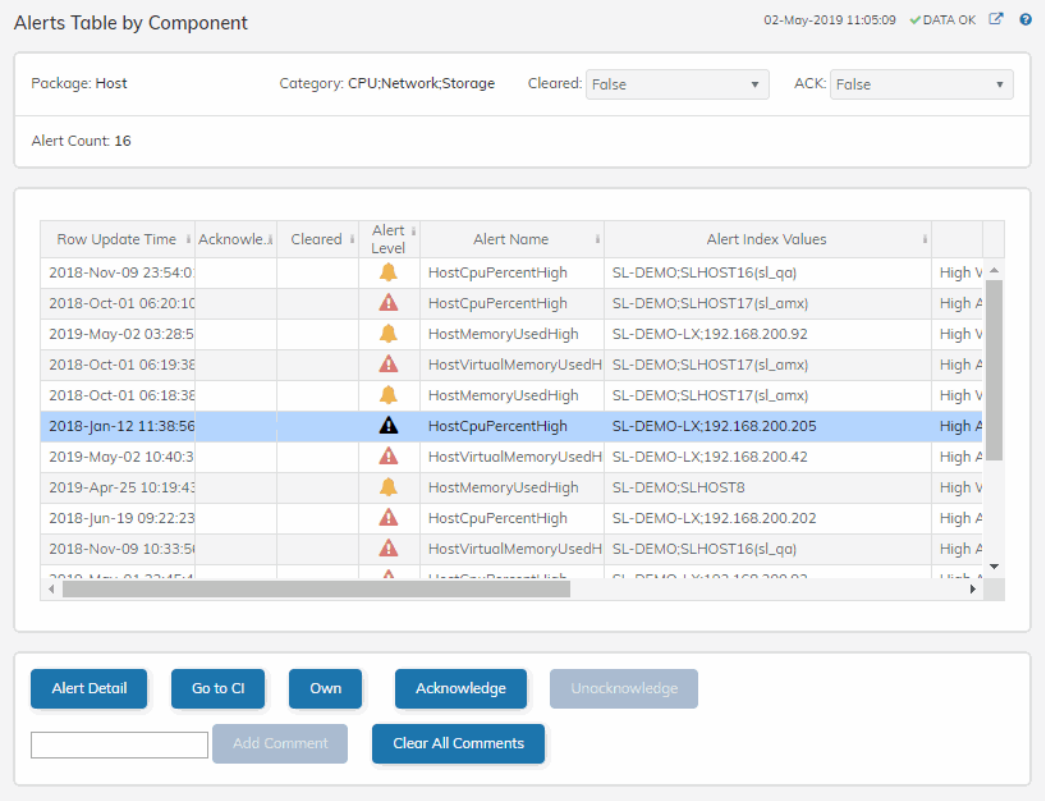

Alert Detail for Component

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.

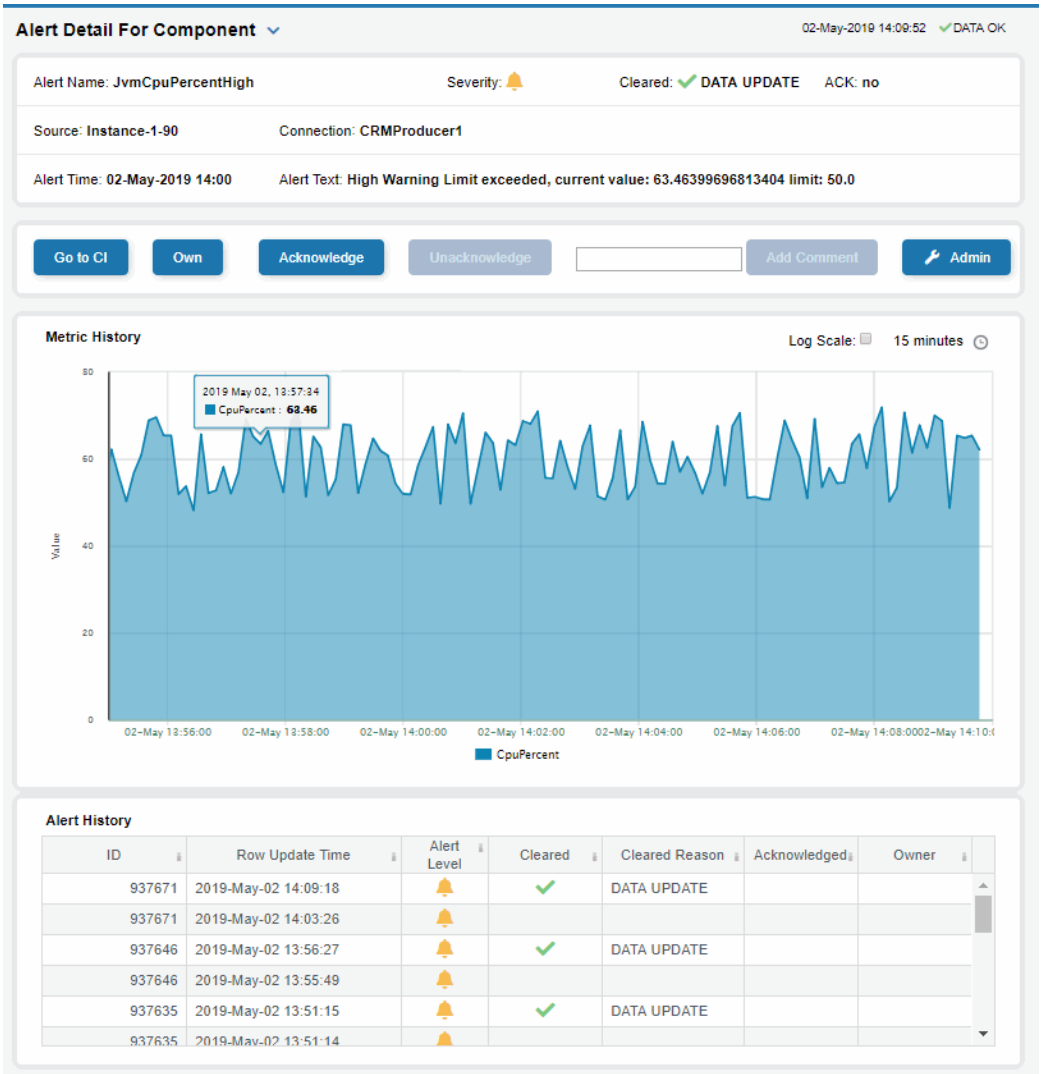

Alerts

These displays present detailed information about all alerts that have occurred in your system. These displays present performance data for your BusinessEvents system. Displays in this View are:

- Alerts Table

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) and **Comment** to add a comment to the alert(s).

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

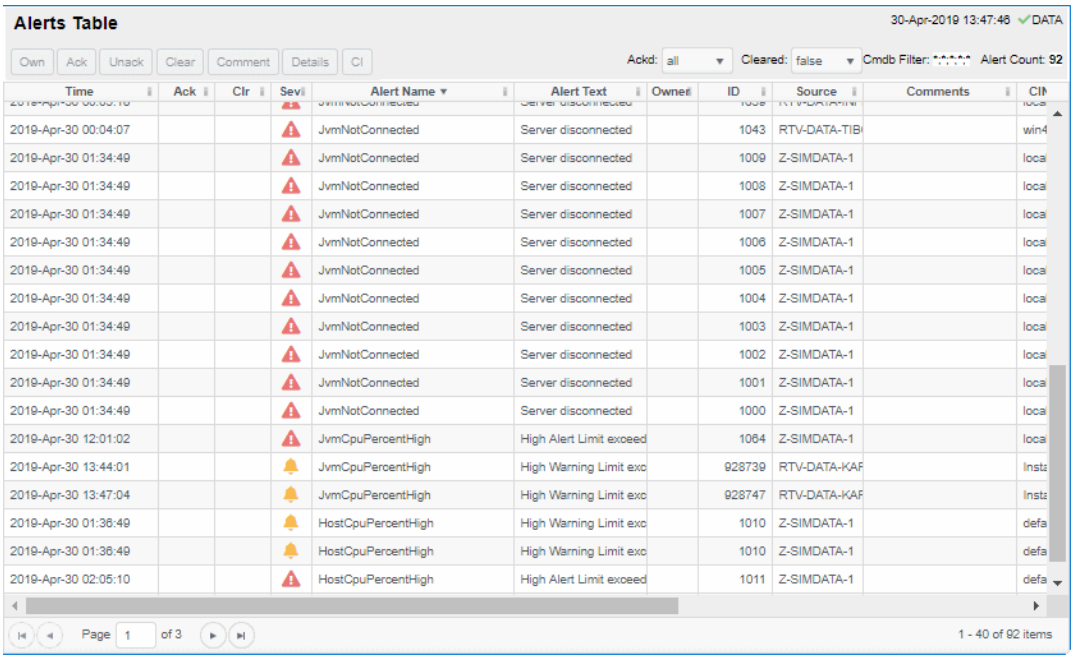

Admin

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**: Opens the Alerts Administration display, which displays active alerts and provides interface to modify, enable and manage alerts.
- **Alert Overrides Admin**: Associated with the Alert Overrides Administration display, which sets and modifies alert overrides. Access to this display is via the **Alert Administration** display. You cannot select this option via the left menu.
- **Alert Engine Status**: Opens the Alert Engine Admin display, which displays a table of available Data Servers and provides an interface to disable and enable the Alert Engine on a Data Server.
- **Cache Table**: Opens the Cache Table display, which allows you to view cached data that RTView is capturing and maintaining, and use this data use this for debugging with SL Technical Support.

Alerts 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.

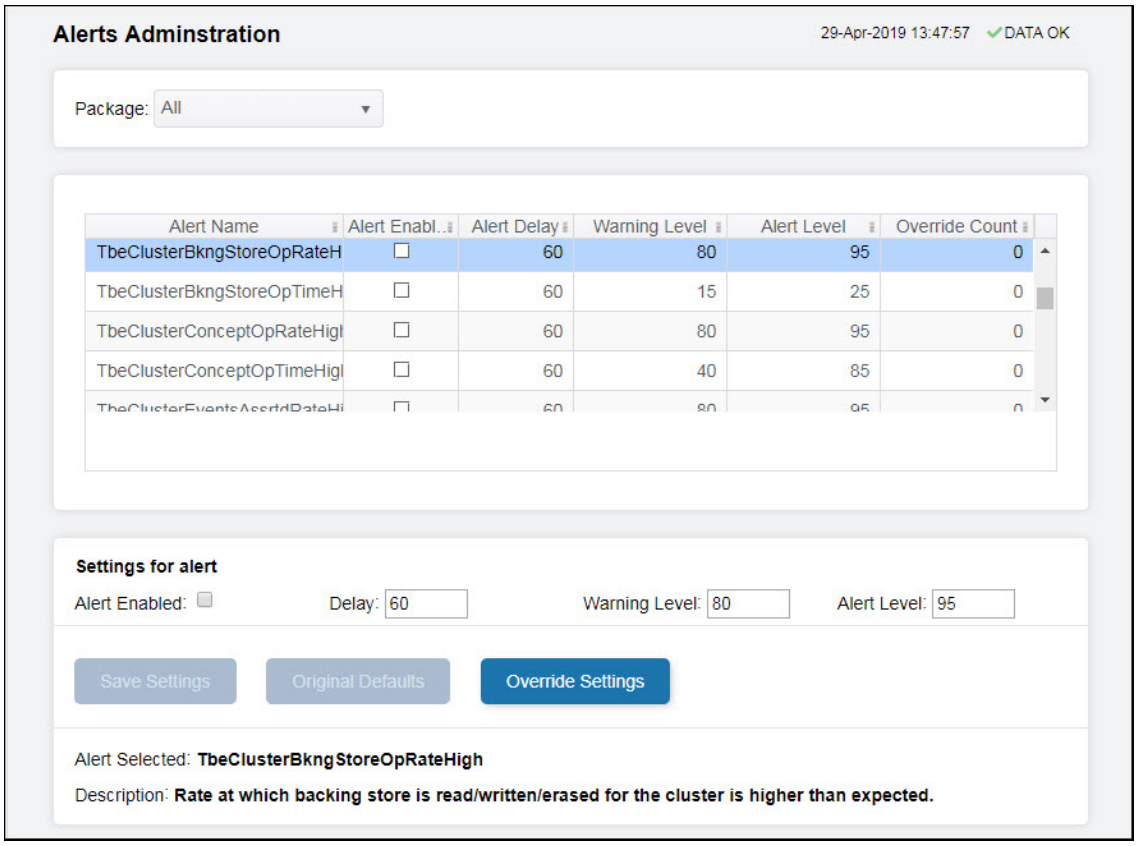

Note: For more information on TIBCO RTView for TIBCO BusinessEvents alerts, see Alert Definitions

Alert Name	The name of the alert.
Alert Enabled	When checked, the alert is enabled globally.
Alert Delay	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. **0** is for immediate execution.
Warning Level	The global warning threshold for the selected alert. When the specified value is exceeded a warning is executed.
Alert Level	The global alarm threshold for the selected alert. When the specified value is exceeded an alarm is executed.
Override Count	The number of times thresholds for this alert have been defined

	individually in the **Tabular Alert Administration** display. A value of: **-0** indicates that no overrides are applied to the alert. **-1** indicates that the alert does not support overrides.

Settings for alert
Select an alert in the table to use the following options:

Alert Enabled	Check / uncheck this box to enable or disable the selected alert globally.
Delay	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.
Warning Level	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.
Alert Level	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. NOTE: For low value-based alerts (such as **EmsQueuesConsumerCountLow**), to set the alarm to occur sooner, increase the Alarm Level value. To set the alarm to occur later, reduce the Alarm Level value.
Save Settings	Click to apply alert settings for the selected alert.
Original Defaults	Click to revert to original alert settings for the selected alert.
Override Settings	Click to set an alert override in the **Alert Overrides Admin** display on the selected alert.

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**.

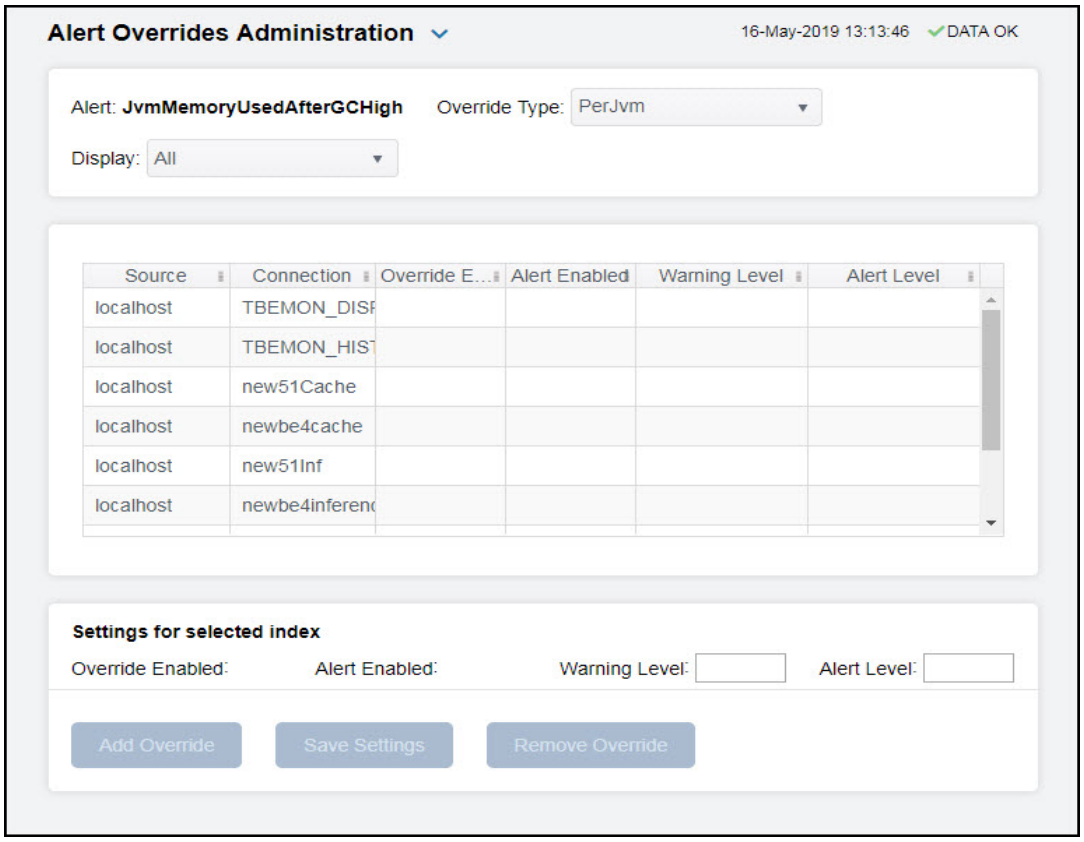

The table lists all the possible overrides that can be defined for the alert you selected from the **Alert Administration** display. Each row in the table represents a different resource or group of resources that can be overridden. When the four last columns are blank, that means the resource has not 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**, **SortAscending**, **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 Engine Admin

This display allows you to enable and disable the alert engine(s) of your Data Server(s) on a per-server basis. This display requires administrator privileges.

Alert Engine Status

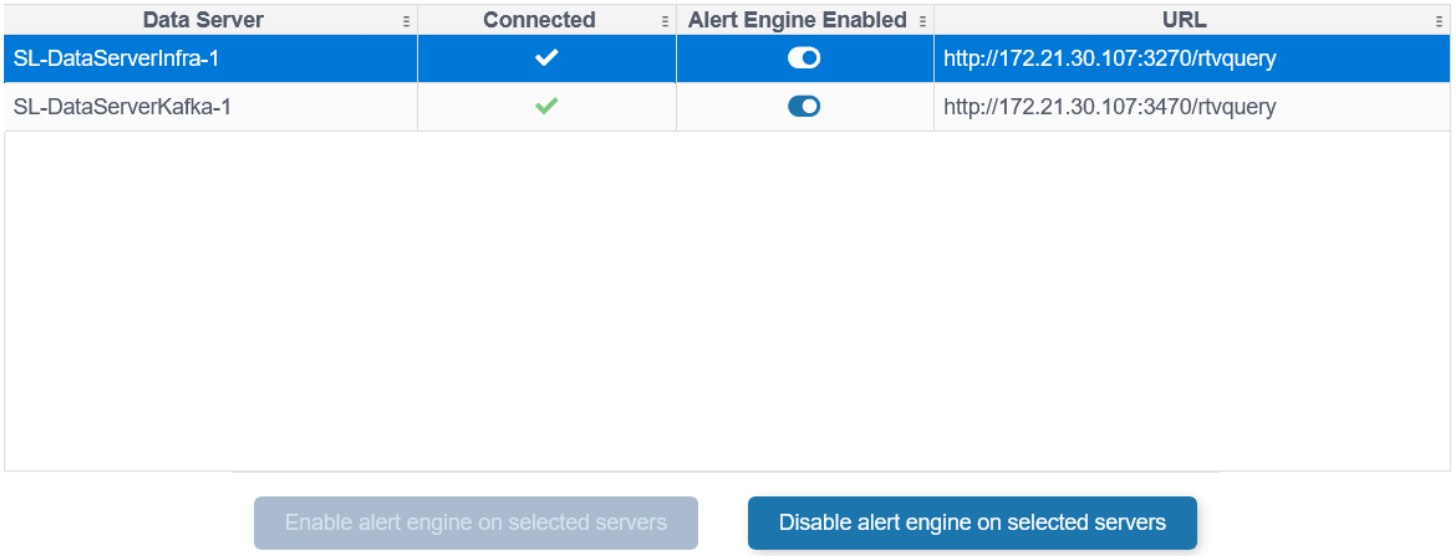

The **Alert Engine Status** table lists Data Servers that are connected to your deployment. The **Connected** column will display whether or not the Data Server is presently connected. If the Data Server is connected, the **Alert Engine Enabled** column will display whether the alert engine for that Data Server is enabled or not.

Disabling the alert engine on a Data Server clears all existing alerts on that server. In the Data Server's RtvAlertTable cache, the "Cleared Reason" column will show MANUAL for each alert that was cleared as a result of disabling the alert engine. No new alerts will be generated by that server until its alert engine is re-enabled..

By default the alert engine is enabled for a Data Server. When a Data Server is restarted, its alert engine is always re-enabled.

Disable Alert Engine

Select one or more Data Servers in the **Alert Engine Status** table, then click

.

A confirmation dialog box will display. Click **OK** to continue and disable the selected alert engine(s).

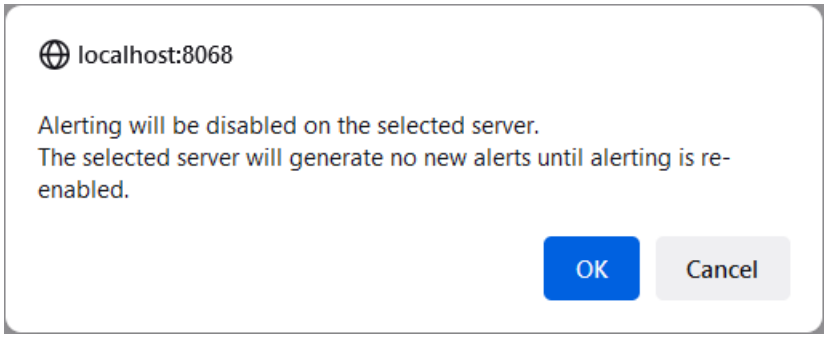

Enable Alert Engine

Select one or more Data Servers in the **Alert Engine Status** table, then click Enable alert engine on selected servers.

A confirmation dialog box will display. Click **OK** to continue and enable the selected alert engine(s).

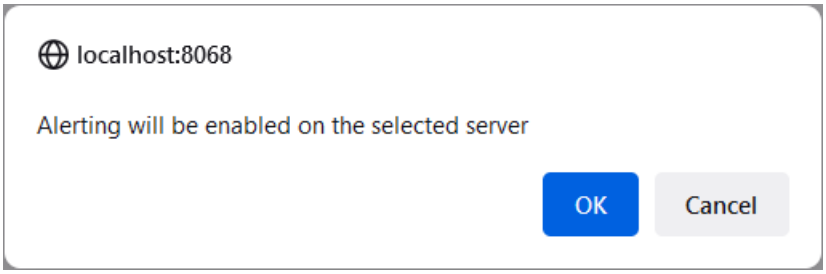

Warning! On Firefox, the enable/disable confirmation dialogs may display a checkbox with the text "Don't allow HOST:8068 to prompt you again". This is from the browser, not from RTView. Do not check that box, as it will prevent the display from working. If it is accidentally selected, clear the Firefox browser cache (**Options -> History -> Clear recent history ...**)

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.

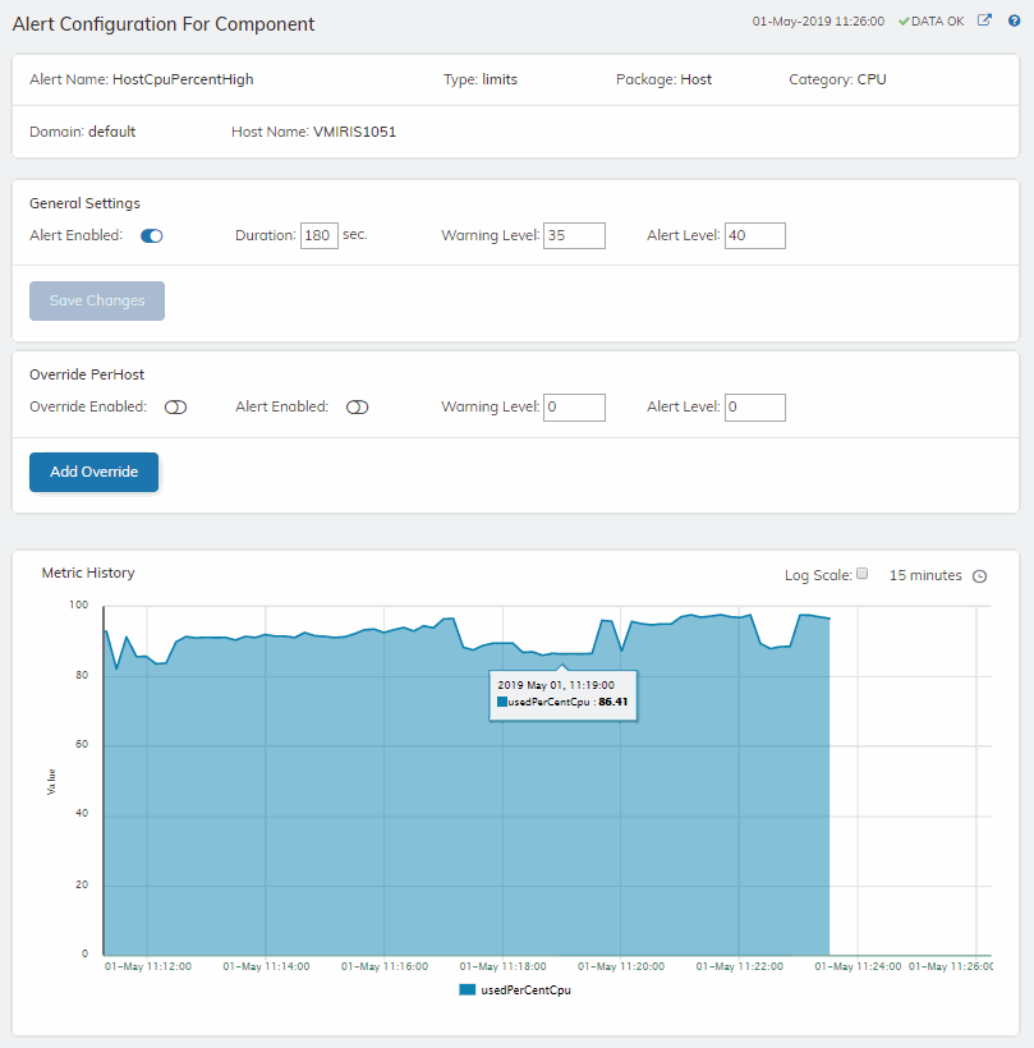

Cache Table

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.

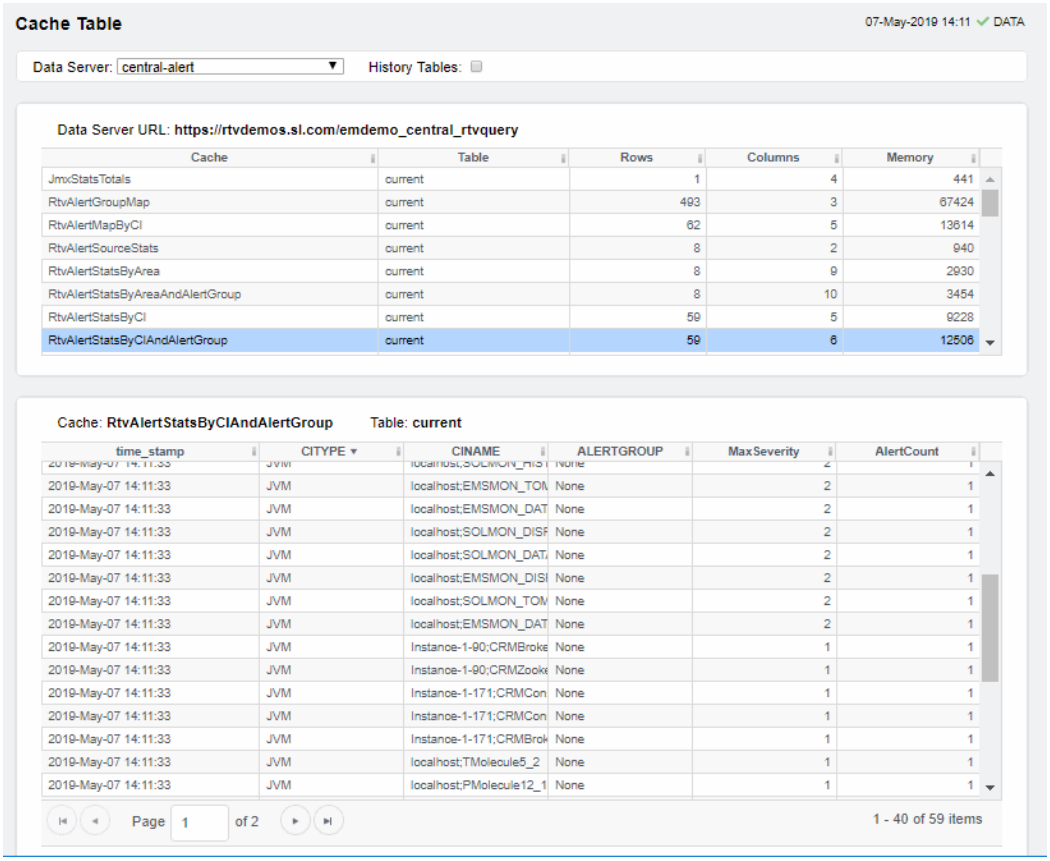

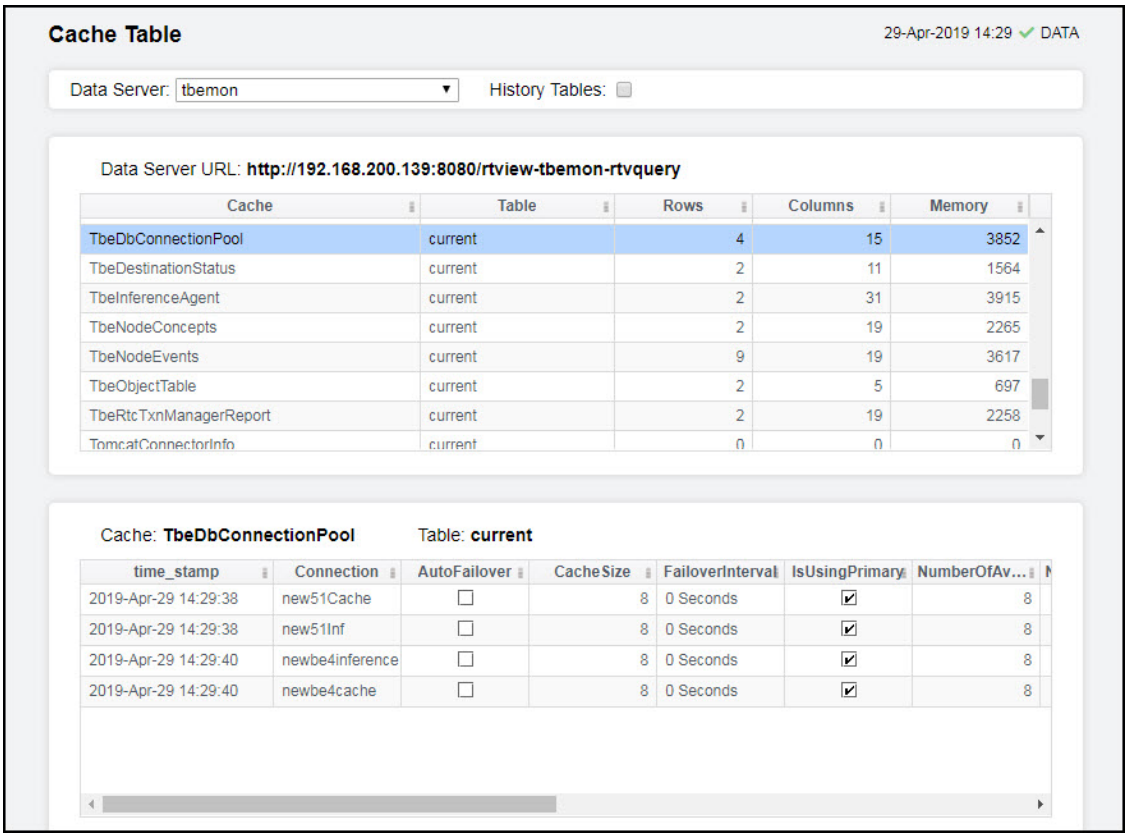
Cache Table
29-Apr-2019 14:29 DATA
Data Server: tbemon
History Tables:
Data Server URL: http://192.168.200.139:8080/rtview-tbemon-rtvquery
Cache
Table
Rows
Columns
Memory
TbeDbConnectionPool current 4 15 3852
TbeDestinationStatus current 2 11 1564
TbeInferenceAgent current 2 31 3915
TbeNodeConcepts current 2 19 2265
TbeNodeEvents current 9 19 3617
TbeObjectTable current 2 5 697
TbeRtcTxnManagerReport current 2 19 2258
TomcatConnectorInfo current 0 0 0
Cache: TbeDbConnectionPool
Table: current
time_stamp
Connection
AutoFailover
CacheSize
FailoverInterval
IsUsingPrimary
NumberOfAv...
2019-Apr-29 14:29:38 new51Cache 8 0 Seconds 8
2019-Apr-29 14:29:38 new51Inf 8 0 Seconds 8
2019-Apr-29 14:29:40 newbe4inference 8 0 Seconds 8
2019-Apr-29 14:29:40 newbe4cache 8 0 Seconds 8

CHAPTER 6 Using the Monitor

This section describes how to read and use Monitor displays. This section includes:

- Overview
- TIBCO BusinessEvents Monitor Views/Displays

Overview

This section describes Monitor displays, how to read Monitor objects, GUI functionality and navigation. This section includes:

- Monitor Main Display
- Trend Graphs
- Time Range
- Tables
- Using the Monitor
- Multiple Windows
- Context Menu
- Export Report
- Rate Calculation in TIBCO BusinessEvents Monitor

Monitor Main Display

The **BE Cluster Table** is the default display of the TIBCO BusinessEvents Monitor. The following figure illustrates the TIBCO BusinessEvents Monitor Clusters display.

Row text and background color change when data expires, or the JMX TCP connection to a given engine breaks. When an alert occurs, the green status LED for the affected engine changes from green to yellow or red, depending on alert severity.

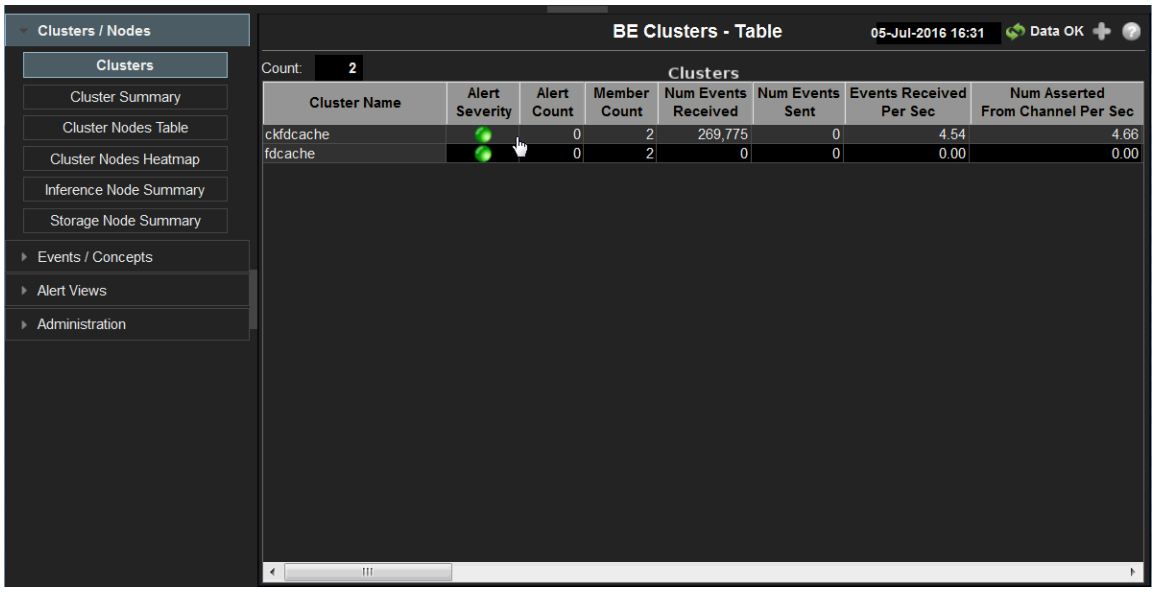

NOTE: It takes about 60 seconds after either a BusinessEvents engine or the Monitor Data Server is started for data to initially appear in Monitor displays. By default, data is collected and displays are refreshed every 30 seconds.

To access online help for any of the Monitor displays select Help (**?**) in the top right corner of the Monitor.

Tables

EMS Monitor tables contain the same data that is shown in the heatmap in the same View. Tables provide you a text and numeric view of the data shown in that heatmap, and additional data not included the heatmap. For example, the **Cluster Nodes Table** display shows the same data as the **Cluster Nodes Heatmap** display.

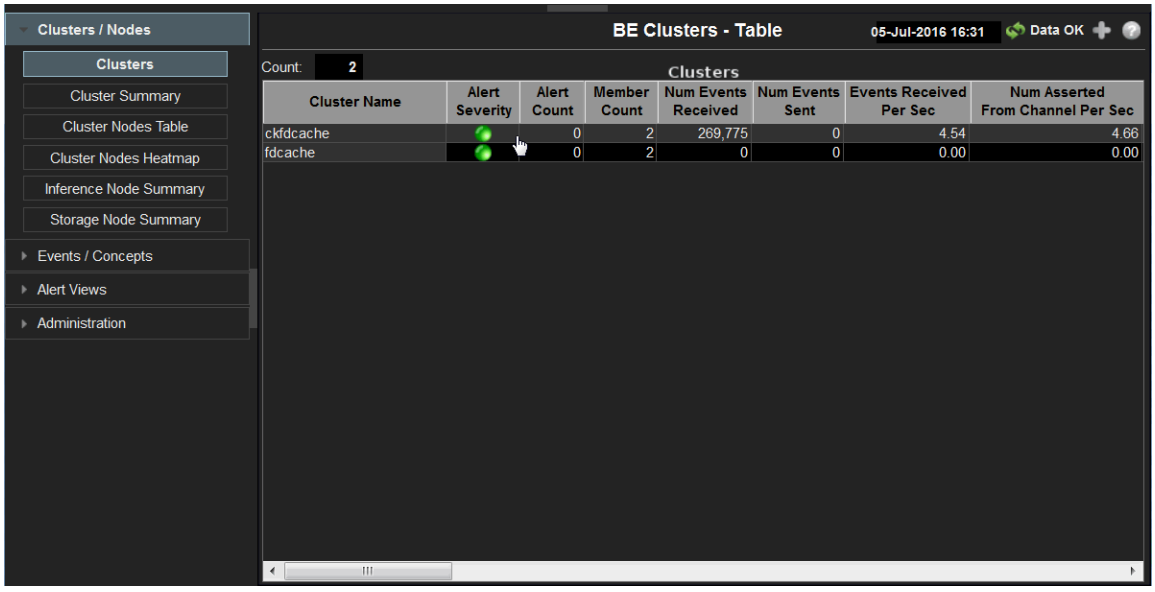

Table rows also sometimes use color to indicate the current most critical alert state for all resources associated with a given row. For example, the color coding is typically as follows:

-- One or more alerts exceeded their critical threshold for one or more associated resources.

--One or more alerts exceeded their warning threshold for one or more associated resources.

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

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

Also see:

- Multiple Column Sorting
- Column Visibility
- Column Filtering
- Column Locking
- Column Reordering
- Saving Settings
- Row Paging
- Row Color Code
- Row Keyboard Selection

Multiple Column Sorting

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

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

The grid's row selection is cleared if the sort is changed or if columns are resized or reordered.

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

Column Visibility

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

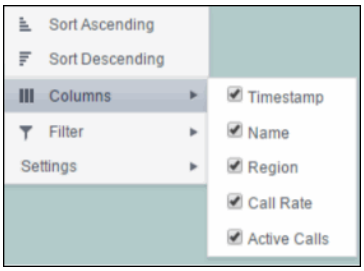

The leftmost column (the row header column) cannot be hidden.

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

Column Filtering

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

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

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

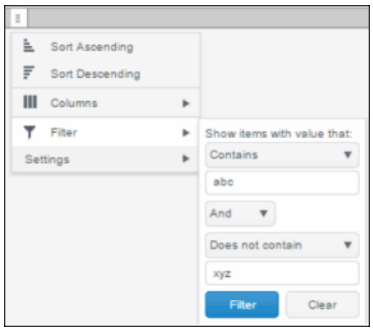

Options in the **Column Filter** dialog vary according to the data type of the selected column:

- **String columns**: You can enter a filter string such as "abc" and, from the dropdown list, select the operator (equal to, not equal to, starts with, contains, etc) to be used when comparing the filter string to each string in the column. All of the filter comparisons on strings are case-insensitive. You can optionally enter a second filter string (e.g. "xyz") and specify if an AND or OR combination should be used to combine the first and second filter results on the column.
- **Numeric columns**: You can enter numeric filter values and select arithmetic comparison operators, (=, !=, >, >=, <, <=). You can optionally enter a second filter value and comparison operator, and specify if an AND or OR combination should be used to combine the first and second filter results.
- **Boolean columns**: You simply select whether matching items should be true or false.

The numeric and boolean filter dialogs are shown below.

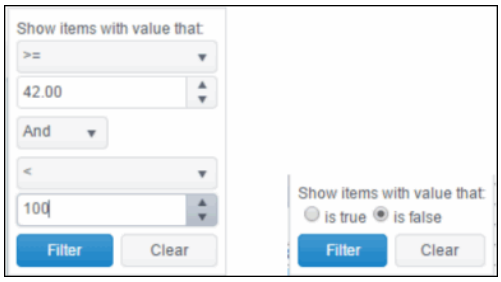

- **Date columns**: You can select a date and time and choose whether matching items should have a timestamp that is the same as, before, or after the filter time. The date is selected by clicking on the calendar icon and picking a date from a calendar dialog. The time is selected by clicking on the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

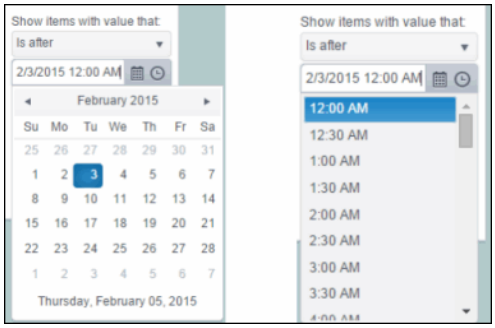

Alternatively, a date and time can be typed into the edit box. The strings shown in a date column are formatted by the Display Server using its time zone. But if a filter is specified on a date column, the date and time for the filter are computed using the client system's time zone. This can be confusing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

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

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

Column Locking

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

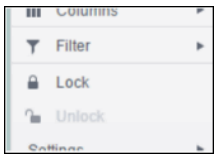

If the row header is enabled, at least one column must remain locked.

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

Column Reordering

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

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

Saving Settings

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

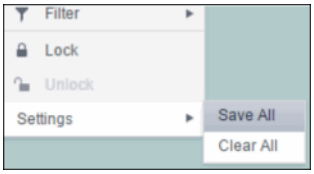

The grid's settings are written as an item in the browser's local storage. The item's value is a string containing the grid's settings. The item uses a unique key comprised of the URL path name, the display name, and the table's RTView object name. If the Thin Client's login feature is enabled, the key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

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

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

Revert Table Settings

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

Row Paging

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

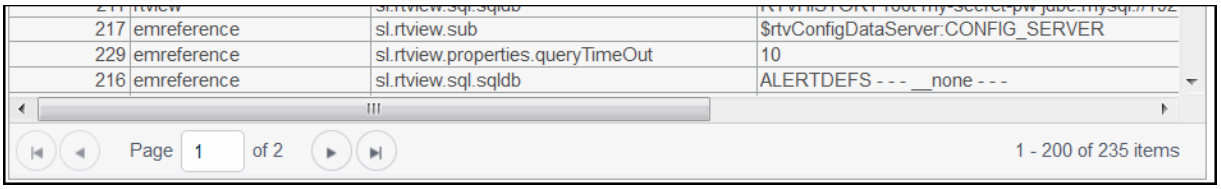

Row Color Code

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

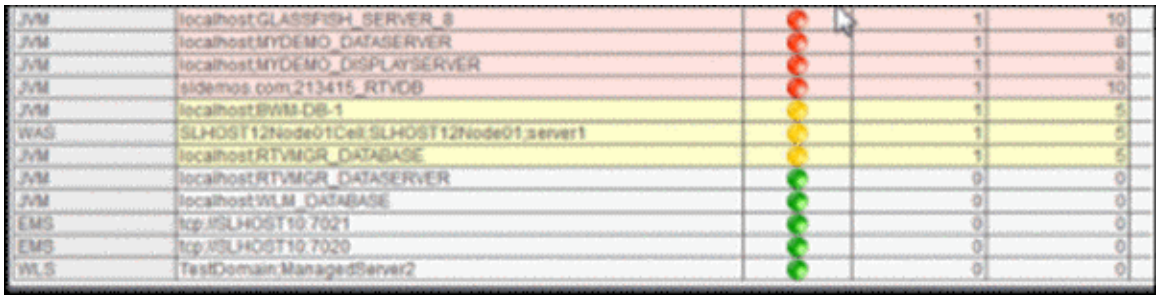

The yellow row color indicates that one or more alerts exceeded their warning threshold for one or more CIs associated with the Service. The red row color indicates that one or more alerts exceeded their critical threshold for the CI associated with the Service (in this case there is a single CI). To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

Row Color Code:

Tables with colored rows indicate the following:

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

Row Keyboard Selection

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

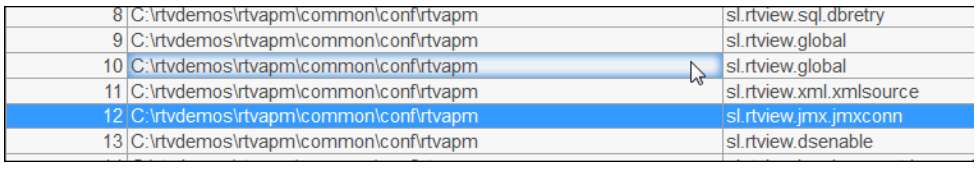

Trend Graphs

TIBCO BusinessEvents Monitor trend graphs enable you to view and compare performance metrics over time. You can use trend graphs to assess utilization and performance trends.

For example, the following figure illustrates a typical TIBCO BusinessEvents Monitor trend graph.

Time Range

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

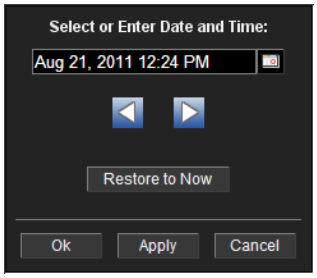

To change the time range click Open Calendar , choose the date and time, then click **OK**. Or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2011 12:24 PM. Click **Apply**. Use the Navigation Arrows to move forward or backward one 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 following figure illustrates mouse-over functionality. In this example, when you mouse-over a single dot, or data point, a pop-up window shows data for that data point.

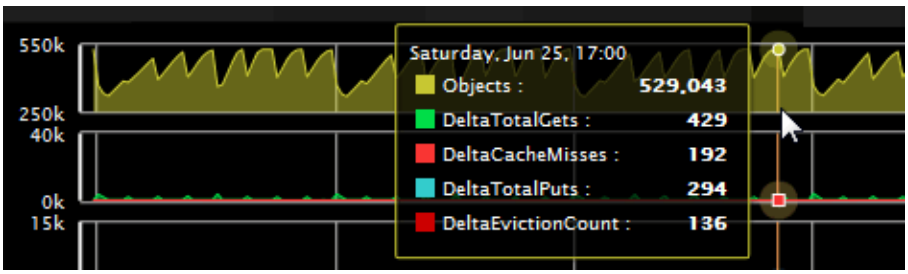

Log Scale

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.

Title Bar Functionality

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

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

	Opens the previous display.
	Opens the display that is up one level.
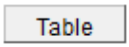	Navigates to a display that is most commonly accessed from the current display. The target display differs among displays.
	Navigates to displays that are most commonly accessed from the current display. The drop-down menu options differ among displays.
	Opens the Alerts Table display in a new window.
	The current date and time. If the time is incorrect, this might indicate that RTView stopped running. When the date and time is correct and the **Data OK** indicator is green, this is a strong indication that the platform is receiving current and valid data.
	The data connection state. Red indicates the data source is disconnected (for example, if the Data Server is not receiving data, or if the Display Server does not receive data from the Data Server, this will be red). Green indicates the data source is connected. When the date and time is correct and the **Data OK** indicator is green, this is a strong indication that the platform is receiving current and valid data.
	The number of items currently in the display.
	Opens an instance of the same display in a new window. Each window operates independently, allowing you to switch views, navigate to other displays in RTView EM, and compare server performance data.

Opens the online help page for the current display.

Multiple Windows

The following illustrates the use of Open New Window in RTView Enterprise.

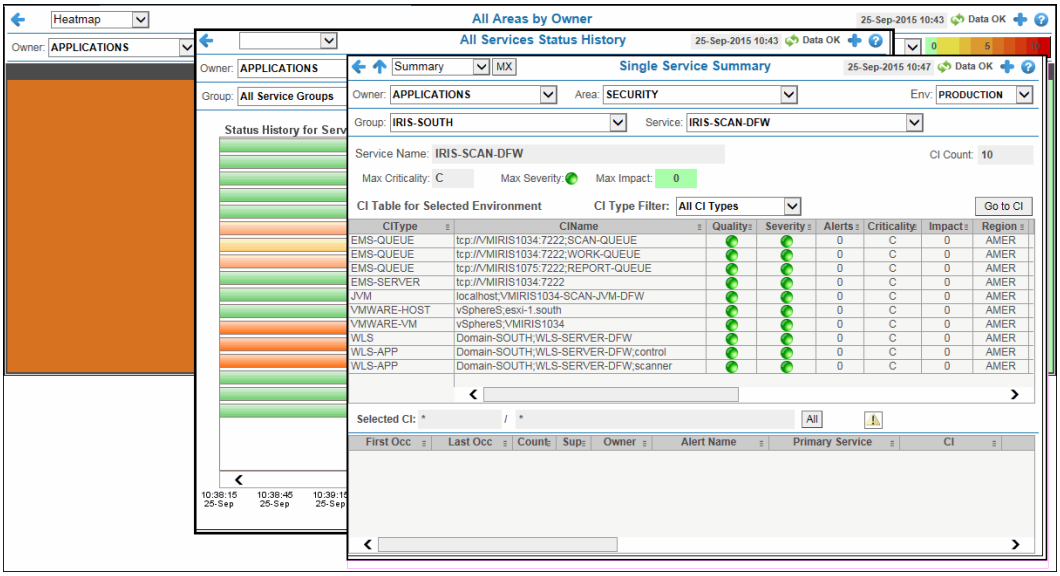

Context Menu

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

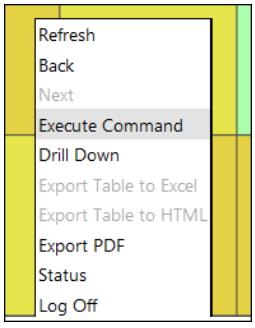

Export Report

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

To generate a report for a display:

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

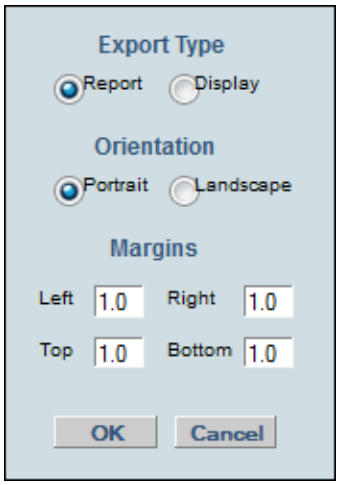

Set the margins and choose the **Export Type**:

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

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

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

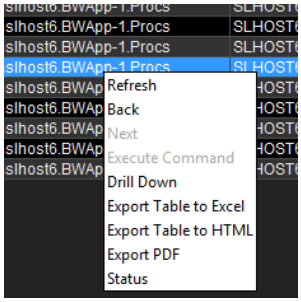

Rate Calculation in TIBCO BusinessEvents Monitor

The majority of rate calculations in BE Monitor use the following property, which is defined in the **\rtvapm\tbemon\conf\rtvapm.tbemon.properties** file:

```
#Query Interval

#

sl.rtview.sub=$tbeQueryIntervalBase:30
```

For **Events Received Per Sec**, for example, if the query interval was defined as 10 (seconds), let's say you had the following:

Query (Time) Interval	Total Number of Events Received
t0	10
t0+10 (seconds)	20
t0+20 (seconds)	30

In the table above, you can see that 10 total events were received during the first query interval. After the second query interval, 20 total events were received. After the third query interval, 30 total events were received. The rate of events received per second for the current query interval is then calculated by taking the total number of events received after the current query interval, subtracting the total number of events received after the previous query interval (to determine the delta), and dividing by the query interval. For example:

Query (Time) Interval	Total Number of Events Received after Query Interval	Delta Time	Delta Number of Events Received (current - previous)	Events Received Per Sec
t0	10	-	-	-
t0+10 (seconds)	20	10	(20-10)	(20-10)/10= 1 Event/sec
t0+20 (seconds)	30	10	(30-20)	(30-20)/10= 1 Event/sec

So, **Events Received Per Sec** is calculated by taking the difference in the number of Events (delta) between the current and previous query interval (30 events - 20 events) and dividing by the current query interval (10 seconds), or:

(30-20)/10 (seconds)= 10/10 = 1 Event per second

TIBCO BusinessEvents Monitor Views/Displays

The TIBCO BusinessEvents Monitor contains the following Views:

- Clusters / Nodes View
- Events / Concepts View
- Alert Views
- Administration

Clusters / Nodes View

These displays present performance data for your BusinessEvents system. Displays in this View are:

- Clusters
- Cluster Summary
- Cluster Nodes Table
- Cluster Nodes Heatmap
- Inference Node Summary
- Storage Node Summary

Clusters

Use this display to check event, concept, and backing store metrics for all of your clusters. Consider keeping this display open to monitor your TIBCO BusinessEvents clusters in general. Each row in the table is a different cluster. Click on a cluster row to view additional cluster details (current and historical) in the **Cluster Summary** **display. The summary display includes trend charts so that you can view key metrics over time.**

Sort the table columns when all the rows cannot fit on the screen. For example, sort the **Alert Status** column so that all nodes with red alerts (🔴) are listed at the top, or sort the **Expired** column so that all expired nodes are listed at the top.

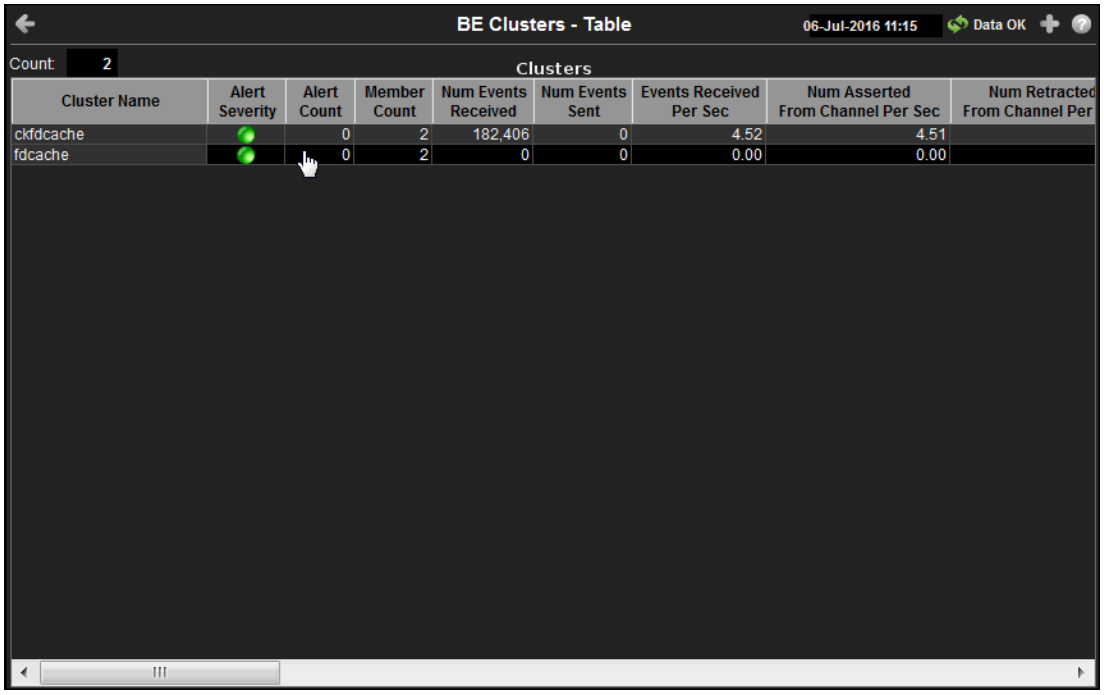

Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu, Table open commonly accessed displays.
- 6,047 The number of items currently in the display.
- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

Clusters Table

Each row in the table is a different cluster, and data in the row columns describe the cluster.

Note: Row columns in this table with an asterisk (*) at the end of the column definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these columns.

Count: The total number of clusters in the table.

Cluster Name	The name of the TIBCO BusinessEvents cluster.
Alert Severity	The severity level of open alerts. Values range from **0** to **2**, where **2** is the greatest Severity: One or more alerts exceeded their ALARM LEVEL threshold. One or more alerts exceeded their WARNING LEVEL threshold. No alert thresholds have been exceeded.
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.
Num Events Received	The total number of events received.*
Num Events Sent	The total number of events sent.*
Events Received Per Sec	The rate of events received in the cluster.
Num Asserted From Channel Per Sec	The rate of events asserted into the Rete network via the channel.
Num Retracted From Channel Per Sec	The rate of events retracted/deleted from the Rete network via the channel.
Num Modified From Channel Per Sec	The rate of events modified in the Rete network via the channel.
Num Rules Fired Rate	The rate of rules fired in the cluster.
Concept Max Get Time	The longest time taken for a "get" operation for any node in the cluster since the cluster was started.*
Concept Max Put Time	The longest time taken for a "put" operation for any node in the cluster since the cluster was started.*
Concept Max Remove Time	The longest time taken for a "remove" operation for any node in the cluster since the cluster was started.*
Concept Max Operation 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/sec	The rate of "get" operations in the cluster.
Concept Puts/sec	The rate of "put" operations in the cluster.
Concept Removes/sec	The rate of "remove" operations in the cluster.
Concept Operations/sec	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 Operation 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/sec	The rate of "erases" in the Backing Store.

Backing Store Loads/sec	The rate of "loads" into the Backing Store.
Backing Store Stores/sec	The rate of "stores" into the Backing Store.
Backing Store Operations/sec	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.
Timestamp	The date and time, relative to the Data Server, that data was last collected for the engine.

Cluster Summary

Use this display to view configuration and utilization data for a single cluster. Select a cluster to view Rete statistics, cache metrics, Backing Store data, and trend data for the cluster.

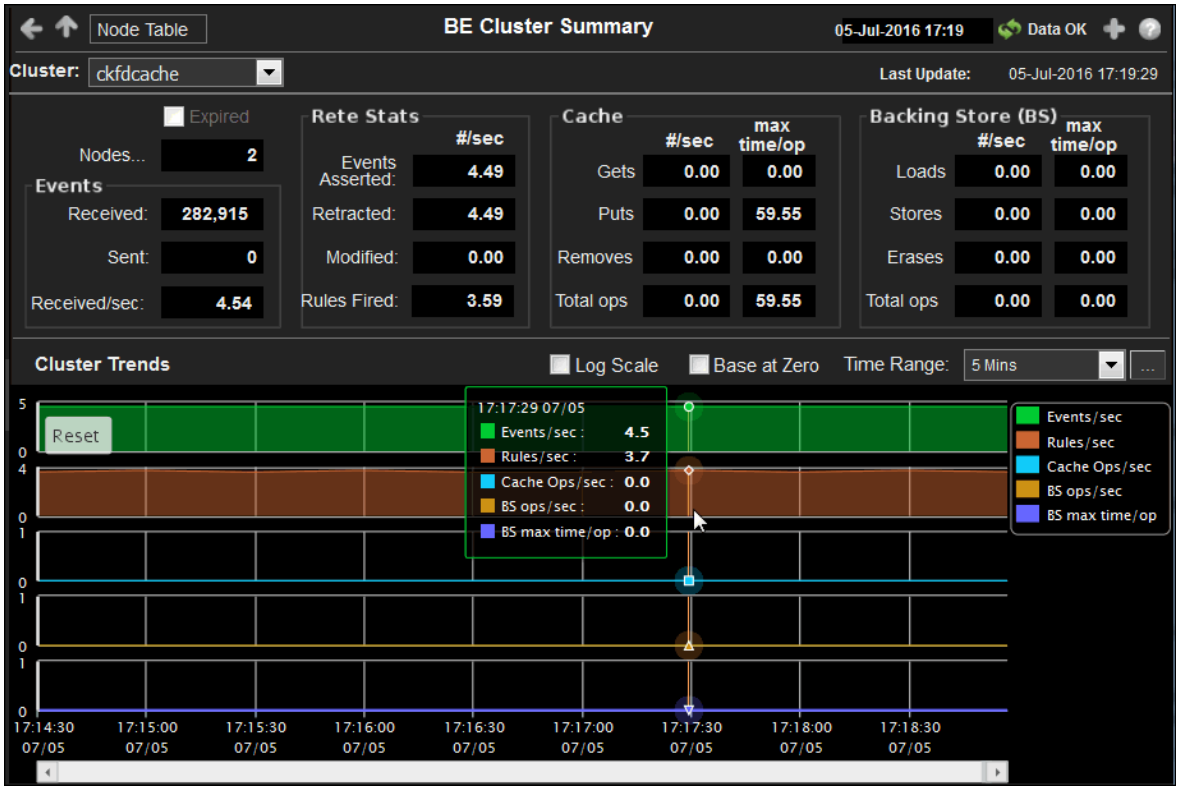

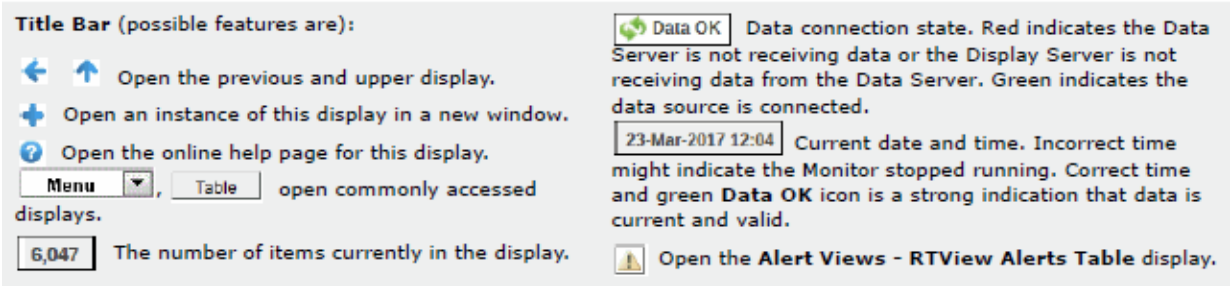

Filter By:

The display might include these filtering options:

Cluster	Choose a cluster for which you want to see metrics.
Last Update	The date and time the data was last updated in the display.

Fields and Data

This display includes:

Note: Fields with an asterisk (*****) at the end of the field definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these fields.

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.
Nodes		Lists the number of nodes on the cluster.
Events	**Received**	The number of events received since the last data update.*****
	Sent	The number of events sent since the last data update.*****
	Received/sec	The rate of events received in the cluster.
Rete Stats	**Events Asserted (#/sec)**	The rate of events asserted into the Rete network.
	Retracted (#/sec)	The rate of events retracted/deleted from the Rete network.
	Modified (#/sec)	The rate of events modified in the Rete network.
	Rules Fired (#/sec)	The rate of rules fired in the Rete network.
Cache	**Gets (#/sec)**	The rate of "get" operations in the L1 cache.
	Gets (max time/op)	The longest time taken for a "get" operation for any node in the cluster since the cluster was started.*****
	Puts (#/sec)	The rate of "put" operations in the cache.
	Puts (max time/op)	The longest time taken for a "put" operation for any node in the cluster since the cluster was started.*****
	Removes (#/sec)	The rate of "removes" in the cache.
	Removes (max time/op)	The longest time taken for a "remove" operation for any node in the cluster since the cluster was started.*****
	Total ops (#/sec)	The rate of operations (gets/puts/removes) in the cache.
	Total ops (max time/op)	The longest time taken for an operation (get/put/remove) for any node in the cluster since the cluster was started.*****
Backing Store	**Loads (#/sec)**	The rate of "load" operations into the backing store.

(BS)

Loads (max time/op)	The longest time taken for a "load" operation in the backing store for any node in the cluster.*
Stores (#/sec)	The rate of "store" operations in the backing store.
Stores (max time/op)	The longest time taken for a "store" operation in the backing store for any node in the cluster.*
Erases (#/sec)	The rate of "erase" operations in the backing store.
Erases (max time/op)	The longest time taken for an "erase" operation in the backing store for any node in the cluster.*
Total ops (#/sec)	The rate of operations (loads/stores/erases) in the backing store.
Total ops (max time/op)	The longest time taken to perform an operation (erase/load/store) in the backing store for any node in the cluster.*

Cluster Trends

Shows the following metrics for the selected cluster.

Events/sec -- 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.

BS ops/sec-- Traces the rate of backstore operations in the cluster.

BS max time/op-- Traces the average maximum time per backstore operation.

Log Scale

This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.

Base at Zero

When this option is checked, zero is set as the Y axis minimum for all graph traces.

Time Range

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

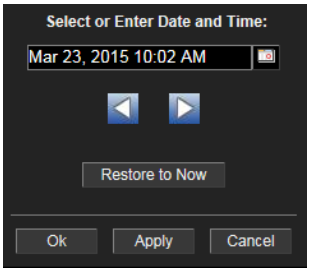

By default, the time range end point is the current time. To change the time range end point, click the button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

Use the navigation arrows to move forward or backward one time period. **Note:** The time period is

determined by your selection from the **Time Range** drop-down menu.

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

Cluster Nodes Table

Use this display to view configuration and utilization data for nodes in a cluster.

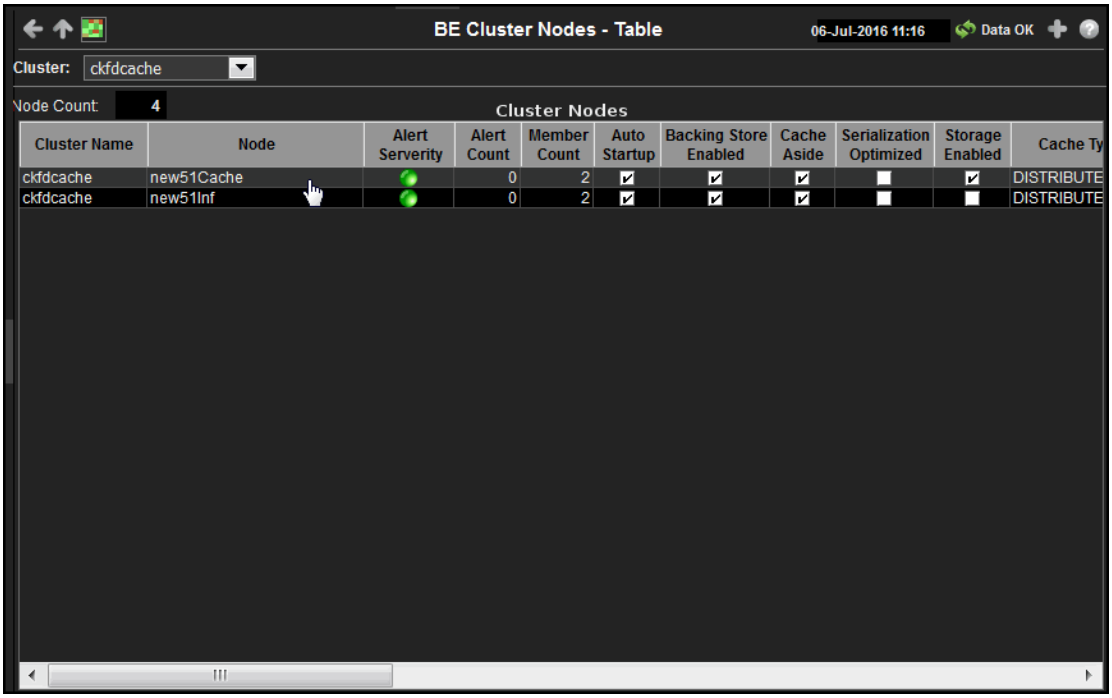

Title Bar (possible features are):

Open the previous and upper display.

Open an instance of this display in a new window.

Open the online help page for this display.

Menu, Table open commonly accessed displays.

6,047 The number of items currently in the display.

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

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

Open the **Alert Views - RTView Alerts Table** display.

Filter By:

The display might include these filtering options:

Cluster Choose a cluster for which you want to 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.

Node Count:	The total number of clusters in the table.
Cluster Name	The name of the TIBCO BusinessEvents cluster.
Node	The name of the node.
Alert Severity	The severity level of open alerts. Values range from **0** to **2**, where **2** is the greatest Severity: One or more alerts exceeded their ALARM LEVEL threshold. One or more alerts exceeded their WARNING LEVEL threshold. No alert thresholds have been exceeded.
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.
Cache Node?	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.
Host	The IP address of the host to which the node is connected.
Port	The port number of the host to which the node is connected.
URL	Uniform Resource Locater, used as an alternative way to specify a JMX connection. When set, the **Host** and **Port** columns are blank (and vice versa).
% CPU Used	The amount of CPU, in percent, used by the node. This value is derived from the java.lang.OperatingSystem MBean.
Heap-Max	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.
Heap-Used	The current amount of memory, in megabytes, in use by the JVM for heap space. This value is provided by standard Java MBeans.
NonHeap 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.
NonHeap 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.
Host OS	The operating system on the host where the node is running.
Connection String	The connection string for the node, which can be the IP address and port of the host that the node is connected to, or the Uniform Resource Locater (which is used as an alternative way to specify a JMX connection).
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.

Timestamp	The date and time, relative to the Data Server, that data was last collected for the node.

Cluster Nodes Heatmap

This display allows you to view utilization data for all nodes in a cluster in a heatmap format. You can view heatmap data for **All Nodes**, **Inference** nodes, or **Cache** nodes by selecting the desired option from the **Nodes** drop down list. When you click on the heatmap for one of the nodes, the detailed data for that particular node displays in the Inference Node Summary display if you selected an inference node, or in the Storage Node Summary display if you selected a cache node.

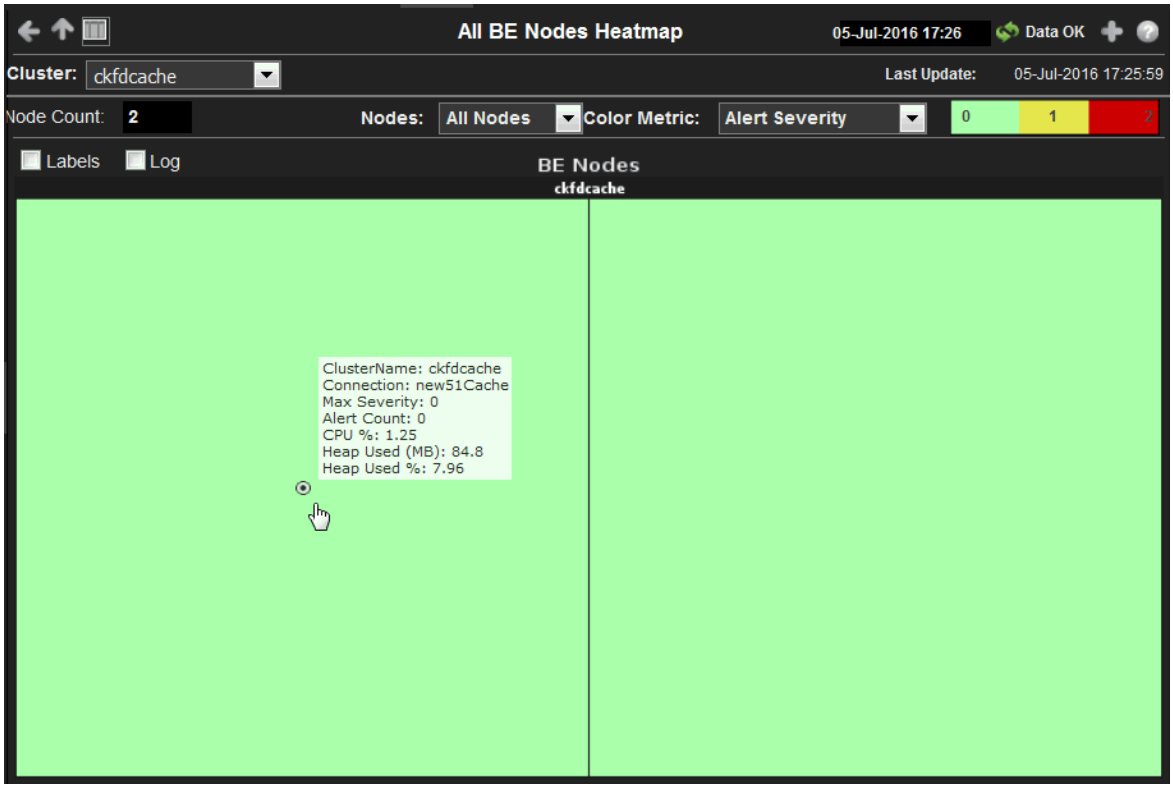

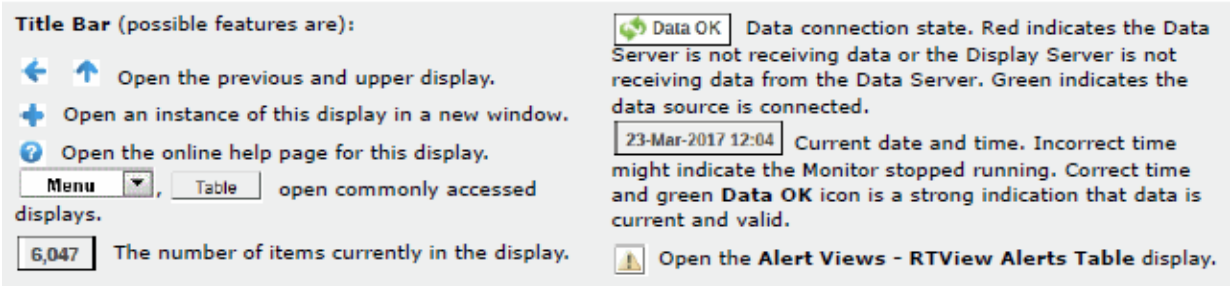

Filter By:

The display might include these filtering options:

Cluster	Choose a cluster for which you want to see metrics.

Last Update The date and time that the display was last updated.

Node Count The total number of nodes in the display.

Nodes Select the type of nodes for which you want to view metrics. You can select from **All Nodes**, **Inference**, and **Cache**. Your selection in this drop down determines the available options in the **Color Metric** drop down.

Labels Select this option to display labels in the heatmap for each of the nodes.

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

Color 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 nodes by cluster, where each rectangle represents a node. Mouse-over any rectangle to display the current values of the metrics for the cluster. Click on a rectangle to drill-down to the associated Storage Node Summary display for a detailed view of metrics for that particular server. The available options in this drop down change depending on your selection in the **Nodes** drop down.

Nodes: All Nodes The following options are available when **All Nodes** is selected from the **Nodes** drop down.

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 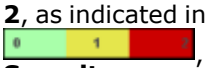 , where **2** is the greatest **Alert Severity**.

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.

JVM % CPU Used The total percentage of JVM CPU used in a given item (index) associated with the rectangle. The color gradient bar 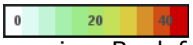 shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of **JvmCpuPercentHigh**, which is **75**. The middle value in the gradient bar indicates the middle value of the range (the default is **38**).

JVM % Memory Used The total percentage of JVM Memory Used in a given item (index) associated with the rectangle. The color gradient bar 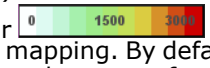 shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold

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.

of **JvmMemoryUsedHigh**, which is **75**. The middle value in the gradient bar indicates the middle value of the range (the default is **38**).

Nodes: Inference

In addition to **Alert Severity**, **Alert Count**, **JVM % CPU Used**, and **JVM % Memory Used**, the following options are also available when **Inference** is selected from the **Nodes** drop down.

Received Events Rate

The rate of events received in a given item (index) associated with the rectangle. The color gradient bar

shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of **TbeNodeChanRecvdRateHigh**, which is **95**. The middle value in the gradient bar indicates the middle value of the range (the default is **48**).

Rules Fired Rate

The rate of rules fired in a given item (index) associated with the rectangle. The color gradient bar

shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of **TbeNodeRuleFiringRateHigh**, which is **95**. The middle value in the gradient bar indicates the middle value of the range (the default is **48**).

Total Rules Fired

The total number of rules fired in a given item (index) associated with the rectangle. The color gradient bar

shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the maximum count of rules fired in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

Nodes: Cache

In addition to Alert Severity, Alert Count, JVM % CPU Used, and JVM % Memory Used, the following options are also available when Cache is selected from the Nodes drop down.

Backing Store Reads/sec

The rate of reads from the backing store in a given item (index) associated with the rectangle. The color gradient bar 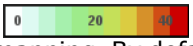

shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of **TbeNodeBackingStoreLoadRateHi**, which is **95**. The middle value in the gradient bar indicates the middle value of the range (the default is **48**).

Backing Store Writes/sec

The rate of writes to the backing store in a given item (index) associated with the rectangle. The color gradient bar

shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of **TbeNodeBackingStoreStoreRateHi**, which is **95**. The middle value in the gradient bar indicates the middle value of the range (the default is **48**).

Backing Store Deletes/sec

The rate of deletes from the backing store in a given item (index) associated with the rectangle. The color gradient bar

shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of **TbeNodeBackingStoreEraseRateHigh**, which is **95**. The middle value in the gradient bar indicates the middle value of the range (the default is **48**).

Concept Gets/sec

The rate of "gets" in a given item (index) associated with the rectangle. The color gradient bar 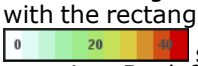

shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of **TbeNodeConceptsGetRateHigh**, which is **95**. The

	middle value in the gradient bar indicates the middle value of the range (the default is **48**).
Concept Puts/sec	The rate of "puts" in a given item (index) associated with the rectangle. The color gradient bar 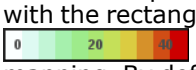 shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of **TbeNodeConceptsPutRateHigh**, which is **95**. The middle value in the gradient bar indicates the middle value of the range (the default is **48**).
Concept Removes/sec	The rater of "removes" in a given item (index) associated with the rectangle. The color gradient bar 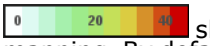 shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of **TbeNodeConceptsRemoveRateHigh**, which is **95**. The middle value in the gradient bar indicates the middle value of the range (the default is **48**).
Object Table Size	The number of objects maintained in the cache in a given item (index) associated with the rectangle. The color gradient bar 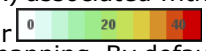 shows the range of the value/color mapping. By default, the numerical values in the gradient bar range from **0** to the alert threshold of **TbeNodeObjectTableSize**, which is **10,000**. The middle value in the gradient bar indicates the middle value of the range (the default is **5,000**).

Inference Node Summary

Use this display to view configuration and utilization data for a single inference node. View a list of all agents on the node, a Run-To-Completion Transaction Manager Report, and trend graphs tracing the rule execution rate for agents on the node. The rule execution rate is relative to the overall CPU and heap utilization for the engine's JVM.

NOTE: An inference node (also known as an engine or processing unit) is the container where one or more inference agents run. Generally, the agents in a given node implement different rule sets, and distributing nodes on different hosts provides fault tolerance and load balancing for the cluster. For details, refer to TIBCO documentation.

Choose a single cluster or **All Clusters** and a node from the drop-down menus.

Change the trend graph **Time Range** to "zoom in" on the graph and see more detail or "zoom out" from the graph to see larger trends over time. To change the time range, click Open Time Range , choose the date and time, and then click **OK**.

Title Bar (possible features are):

Open the previous and upper display.

Open an instance of this display in a new window.

Open the online help page for this display.

Menu, Table open commonly accessed displays.

6,047 The number of items currently in the display.

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

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

Open the **Alert Views - RTView Alerts Table** display.

Filter By:

The display might include these filtering options:

Cluster	Choose a cluster for which you want to view metrics.
Node	Choose a node for which you want to view metrics.

Fields and Data:

Last Update	The date and time the data in the display was last updated.
Cluster Name:	The name of the TIBCO BusinessEvents cluster with which the node is a member.
BE Version:	The approximate TIBCO BusinessEvents version, as configured by the **connection** property. The exact version information is not available via JMX.
Node ID:	A unique string that identifies the node.
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.
% CPU:	The percent of CPU used by the engine process. This value is provided by standard Java MBeans.
Heap used:	The current amount of memory, in megabytes, in use by the JVM for heap space. This value is provided by standard Java MBeans.

Heap max	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.
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.
Auto Startup	When checked (**true**), this feature is enabled.
Cache Aside	When checked (**true**), this feature is enabled.
Backing Store Enabled	When checked (**true**), this feature is enabled.
Storage Enabled	When checked (**true**), this feature is enabled.
Serialization Optimized	When checked (**true**), this feature is enabled.

RTC TXN Manager Report

Note: Fields in this display with an asterisk (*****) at the end of the field definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these fields.

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.
Avg Action Txn Millisec	The average amount of time taken for an action transaction, in milliseconds.*****
Avg Cache Queue Wait Time Millisec	The average cache queue wait time, in milliseconds.*****
Avg Cache Txn Millisec	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 Millisec	The average database queue wait time, in milliseconds.*****
Avg DB Txn Millisec	The average amount of time taken for a database transaction, in milliseconds.*****
Avg Successful Txn Time Millisec	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	The total number of database transactions that have been completed.*****

Completed	
Total Errors	The total number of errors.*
Total Successful Txns	The total number of successful transactions.*

Agents for this Node Table

Each row in the table is an agent associated with the node, with data in the row columns describing the agent.

Note: Row columns in this table with an asterisk (*) at the end of the column definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these columns.

Agent Count:	The number of agents currently in the table.
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/sec	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.

Timestamp The date and time, relative to the Data Server, that data was last collected for the agent.

Trend Graph Shows metrics for the selected node.

- **% CPU** -- Traces the amount of CPU used, in percent, by the node.
- **Rules/sec --** Traces the number of rules processed, per second, by the agent.
- **Heap-max** -- Traces the maximum amount of heap space, in bytes, used by the node since the agent was started.
- **Heap-used** -- Traces the current amount of heap space, in bytes, used by the agent.

Rules/sec for Agent Choose an agent from the drop-down menu.

Log Scale This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.

Base at Zero When this option is checked, zero is set as the Y axis minimum for all graph traces.

Time Range

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

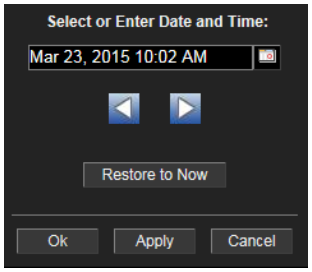

By default, the time range end point is the current time.

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

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

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

Storage Node Summary

Use this display to view configuration details for a single cache node, the database connection pool status, as well as a list of all caches that are backed by the backing store (database). Also view trend graphs that trace utilization metrics such as CPU and heap memory usage.

NOTE: A storage node (also known as a cache node) provides fast access to events and concepts required during each RTC by the inference engines. Storage nodes also serve as buffers for reads and writes between the cluster and the backing store. For details, refer to TIBCO documentation.

Choose a single cluster or **All Clusters** and a node from the drop-down menus.

Change the trend graph **Time Range** to "zoom in" on the graph and see more detail or "zoom out" from the graph to see larger trends over time. To change the time range click Open Time Range ..., choose the date and time, then click **OK**.

Title Bar (possible features are):

- Open the previous and upper display.
- Open an instance of this display in a new window.
- Open the online help page for this display.
- Menu, Table open commonly accessed displays.
- 6,047 The number of items currently in the display.
- Data OK Data connection state. Red indicates the Data Server is not receiving data or the Display Server is not receiving data from the Data Server. Green indicates the data source is connected.
- 23-Mar-2017 12:04 Current date and time. Incorrect time might indicate the Monitor stopped running. Correct time and green **Data OK** icon is a strong indication that data is current and valid.
- Open the **Alert Views - RTView Alerts Table** display.

Filter By:

The display might include these filtering options:

Cluster: Choose a cluster to see metrics for.

Node:	Choose a node to see metrics for.

Fields and Data

Last Update	The date and time the data was last updated in the display.
Cluster Name:	The name of the TIBCO BusinessEvents cluster with which the node is a member.
BE Version:	The approximate TIBCO BusinessEvents version, as configured by the **connection** property. The exact version information is not available via JMX.
Node ID:	A unique string that identifies the node.
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.
% CPU:	The amount of CPU, in percent, used by the node. This value is provided by standard Java MBeans.
Heap used:	The current amount of memory, in megabytes, in use by the JVM for heap space. This value is provided by standard Java MBeans.
Heap Max:	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.
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.
Auto Startup	When checked (**true**), this feature is enabled.
Cache Aside	When checked (**true**), this feature is enabled.
Backing Store Enabled	When checked (**true**), this feature is enabled.
Storage Enabled	When checked (**true**), this feature is enabled.
Serialization Optimized	When checked (**true**), this feature is enabled.

DB Connection Pool

Values describe status of the pool of database connections used by the cache agent to move data between the local caches and the database.

Note: Fields in this region with an asterisk (*****) at the end of the field definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these fields.

Pool State	The state of the database connection pool.*****
Auto Failover	The number of times auto failover has occurred.*****
Failover Interval	The number of seconds taken for failover to take place.*****
Cache Size	The cache size.*****
# Connections Available	The total number of connections available.*****
# Connections in Use	The total number of connections currently in use.*****

Backing StoreTable

A cache node manages access to current events and concepts, buffering as necessary between local memory and a database. The Backing Store table provides a list of caches and the database select/insert/delete statistics for each cache.

Note: Row columns in this table with an asterisk (*****) at the end of the column definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these columns.

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/sec	The rate of "delete" operations in the node.
Loads/sec	The rate of "load" operations in the node.
Stores/sec	The rate of "store" operations in the node.

Object Table Trends

Shows metrics for the selected cluster/node combination:

% CPU -- Traces the amount of CPU used, in percent, by the engine.

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.

Max Heap (MB)-- Traces the maximum amount of memory, in megabytes, that can be used by the JVM for heap space.

Heap (MB) -- Traces the current heap space, in megabytes, in use by the JVM.

Log Scale

This option should be used when the range of your data is very broad. When checked, the values are displayed using a logarithmic scale rather than using the actual values so that data on the extreme ends of the scale can be viewed more effectively. For example, if you have data that ranges from the tens to the thousands, the data in the range of the tens will be neglected visually if you do not check this option.

Base at Zero

When this option is checked, zero is set as the Y axis minimum for all graph traces.

Time Range

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

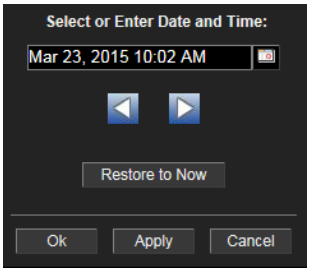

By default, the time range end point is the current time.

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

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

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

Events / Concepts View

These displays present performance data for your BusinessEvents system. Displays in this View are:

- Agent Events
- Agent Event Summary
- Event Cache Hits
- Event Hit Summary
- Concept Cache Hits
- Concept Hit Summary
- Channels
- All Inference Agents
- All RTC Reports

Agent Events

View run-time statistics for a selected group of agents. With TIBCO BusinessEvents, events are cached when they are out of scope, and deleted or persisted to the backing store when they are no longer useful. Clicking on a row in the table displays access patterns over time for the event in the Agent Event Summary display.

NOTE: Events cause rules to execute in the BusinessEvents Rete network. Events can be created by external phenomena, such as the arrival of a JMS message, or internally when rules are processed. When an event enters the Rete network, it causes a run-to-completion cycle which continues until no further rules can be processed. Each named event that can be handled by a BusinessEvents application is specified at build time in BusinessEvents studio. For details, refer to TIBCO documentation.

Sort the table columns when all the rows cannot fit on the screen. For example, sort the **Expired** column so that all expired nodes are listed at the top.

Title Bar (possible features are):

Open the previous and upper display.

Open an instance of this display in a new window.

Open the online help page for this display.

Menu, Table open commonly accessed displays.

6,047 The number of items currently in the display.

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

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

Open the **Alert Views - RTView Alerts Table** display.

Filter By:

The display might include these filtering options:

Cluster:	Select the cluster containing the node and agent for which you want to view metrics.
Node:	Select a node containing the agent for which you want to view metrics.
Agent	Select the agent for which you want to view metrics.

Fields and Data:

Last Update:	The date and time the data on the display was last updated.

Agent Events Table:

Each row in the table is a different event. Data in the row columns describe the event. The following fields are added by Monitor collection. The assertions/sec, modified/sec, and retracted/sec metrics are calculated from the corresponding counters as the delta between two successive samples divided by the polling interval.

Note: Row columns in this table with an asterisk (*****) at the end of the column definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these columns.

Event Count:	The total number of events in the table.
Event	The name of the event.
Avg Time Pre RTC	The average amount of time taken for the event to begin its run to completion cycle.*****

Avg Time in RTC	The average amount of time taken for the event to complete (once it has started) its run to completion cycle.*****
Avg Time Post RTC	The average amount of time taken by the event after its run to completion cycle has ended.*****
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/sec	The rate of L1 cache hits.
L1 Cache Misses/sec	The rate of L1 cache misses.
# Recovered /sec	The rate of recovered data.
Assertions/sec (Agent)	The rate of event assertions into the Rete network by the agent.
Assertions/sec (Channel)	The rate of event assertions into the Rete network via the channel.
Modifies/sec (Agent)	The rate of event modifications in the Rete network by the agent.
Modifies/sec (Channel)	The rate of event modifications in the Rete network via the channel.
Retractions/sec (Agent)	The rate of event retractions/deletions from the Rete network by the agent.
Retractions/sec (Channel)	The rate of event retractions/deletions from the Rete network via the channel.
Expired	When checked, performance data has not been received within the time specified (in seconds) in the **Expire Time** field in the **Duration** region in the RTView Configuration Application > (**Project Name**) > **Solution Package Configuration** > **TIBCO BusinessEvents** > **DATA STORAGE** tab. The **Delete Time** field (also in the **Duration** region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
Timestamp	The date and time, relative to the Data Server, that data was last collected for the engine.

Agent Event Summary

View detailed performance metrics for an agent's event. You can view cache, RTC, event statistics by channel, and event trend data over a specified period of time.

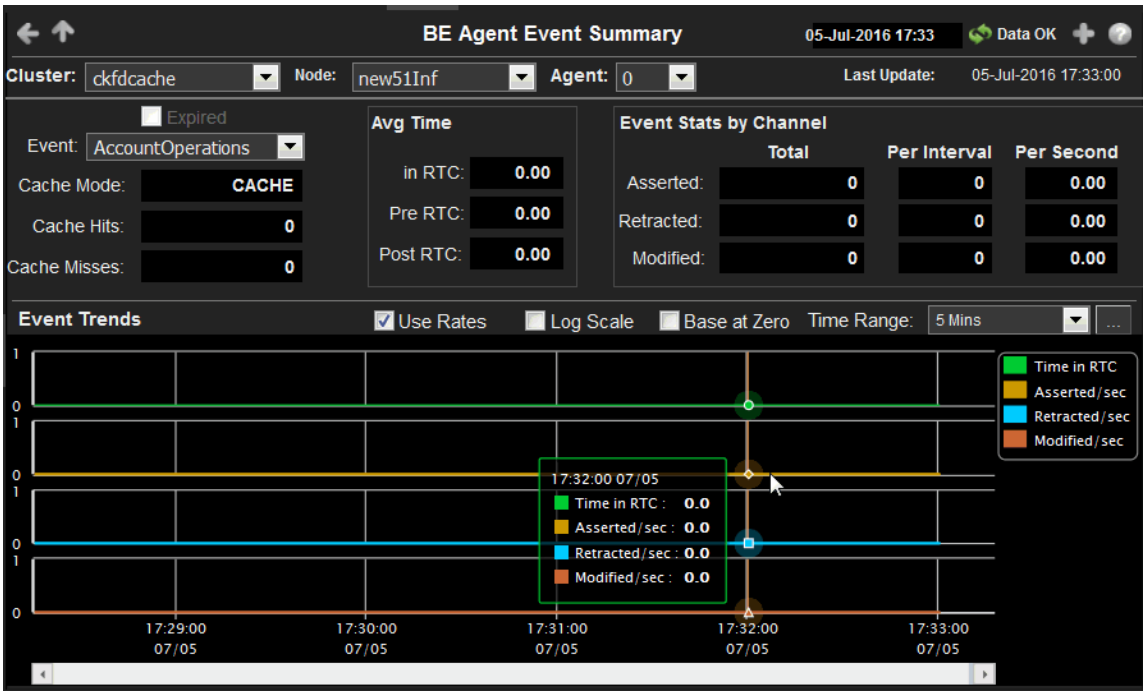

Title Bar (possible features are):

Open the previous and upper display.

Open an instance of this display in a new window.

Open the online help page for this display.

Menu, Table open commonly accessed displays.

6,047 The number of items currently in the display.

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

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

Open the **Alert Views - RTView Alerts Table** display.

Filter By:

The display might include these filtering options:

Note: Fields in this display with an asterisk (*) at the end of the column definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these fields.

Cluster:	Select the cluster for which you want to see metrics.
Node:	Select the node for which you want to see metrics.
Agent	Select the agent for which you want to see metrics.
Last Update	The date and time in which the data was last updated.
Expired	When checked (**true**), the Monitor has not received a response from the event for the amount of time specified by the **$tbeRowExpirationTime** property (the default is **120** seconds). When the amount of time specified by the **$tbeRowExpirationTimeForDelete** property elapses (the default is one day), the event data is deleted from the cache and display.
Event	The name of the event.
Cache Mode	Lists the mode used by the event, which can be either **CACHE** (only) or **MEMORY** (only).*
Cache Hits	The number of times data has been searched for in the L1 cache since the last data update.*
Cache	The number of times data has been searched for in the L1 cache, but was not found,

Misses	since the last data update.*	
Avg Time	**in RTC**	The average amount of time taken for the event to complete (once it has started) its run to completion cycle.*
	Pre RTC	The average amount of time taken for the event to begin its run to completion cycle.*
	Post RTC	The average amount of time taken by the event after its run to completion cycle has ended.*
Event Stats by Channel	**Asserted Total**	The total number of times the event was asserted into the Rete network via the channel.*
	Asserted Per Interval	The number of times the event was asserted into the Rete network via the channel since the last data update.*
	Asserted Per Second	The rate of event assertions into the Rete network via the channel.
	Retracted Total	The total number of times the event was retracted/deleted from the Rete network via the channel.*
	Retracted Per Interval	The number of event retractions/deletions from the Rete network.
	Retracted Per Second	The rate of event retractions/deletions from the Rete network via the channel.
	Modified Total	The total number of times the event was modified in the Rete network via the channel.*
	Modified Per Interval	The number of event modifications in the Rete network via the channel.
	Modified Per Second	The rate of event modifications in the Rete network via the channel.
Expired	When checked, performance data has not been received within the time specified (in seconds) in the **Expire Time** field in the **Duration** region in the RTView Configuration Application > (**Project Name**) > **Solution Package Configuration** > **TIBCO BusinessEvents** > **DATA STORAGE** tab. The **Delete Time** field (also in the **Duration** region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.	
Event Trends	Shows metrics for the selected event: **Time in RTC**-- Traces the event spends in the run to completion cycle. **Asserted(/sec)**-- Traces the number of events asserted into the Rete network (or the rate of event assertions per second depending on **Use Rates** setting). **Retracted(/sec)**-- Traces the number events retracted from the Rete network (or rate of event retractions per second depending on **Use Rates** setting). **Modified(/sec**)-- Traces the number of events modified in the Rete network (or rate of events modified per second depending on **Use Rates** setting).	
	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.
	Base at Zero	When this option is checked, zero is set as the Y axis minimum for all graph traces.

Time Range

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

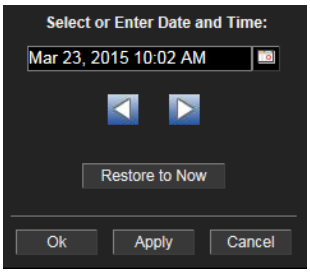

By default, the time range end point is the current time. To change the time range end point, click the button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

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

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

Event Cache Hits

View cache performance metrics per event for a single cluster or **All Clusters**.

NOTE: Events cause rules to execute in the BusinessEvents Rete network. Events can be created by external phenomena, such as the arrival of a JMS message, or internally when rules are processed. When an event enters the Rete network, it causes a run-to-completion cycle which continues until no further rules can be processed. Each named event that can be handled by a BusinessEvents application is specified at build time in BusinessEvents studio. For details, refer to TIBCO documentation.

Choose a single cluster or **All Clusters** and a node from the drop-down menus.

Sort the table columns when all the rows cannot fit on the screen. For example, sort the **Expired** column so that all expired nodes are listed at the top.

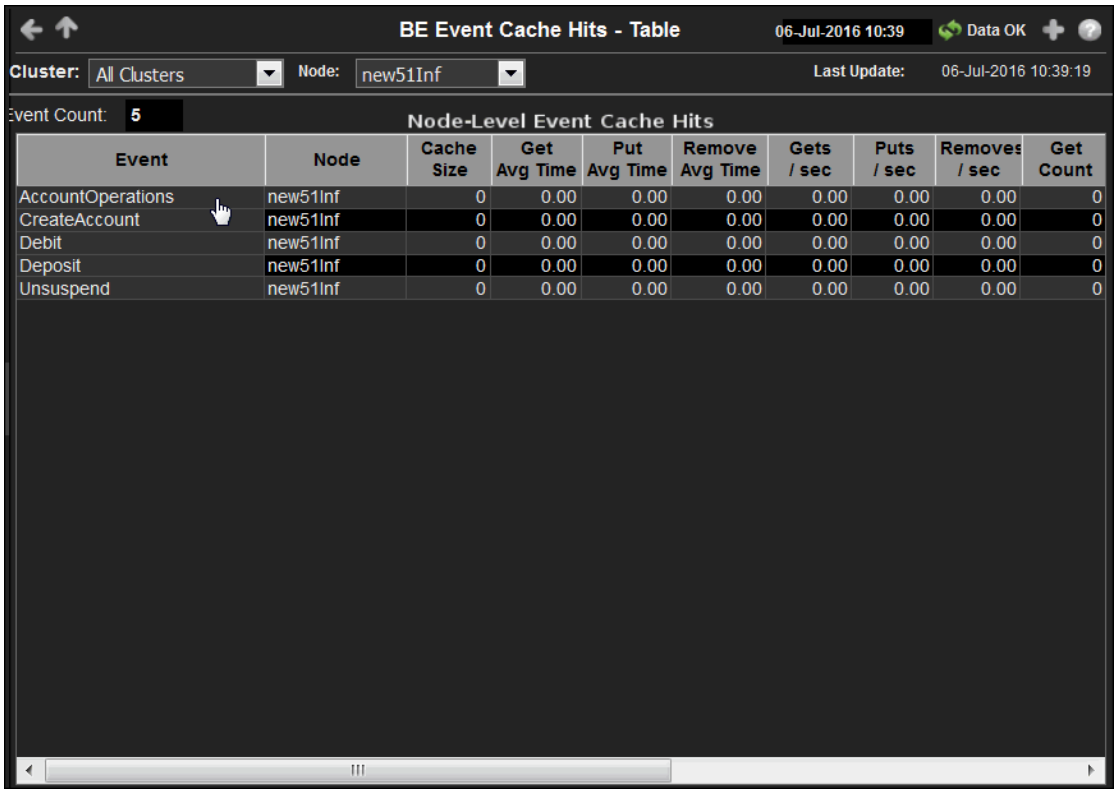

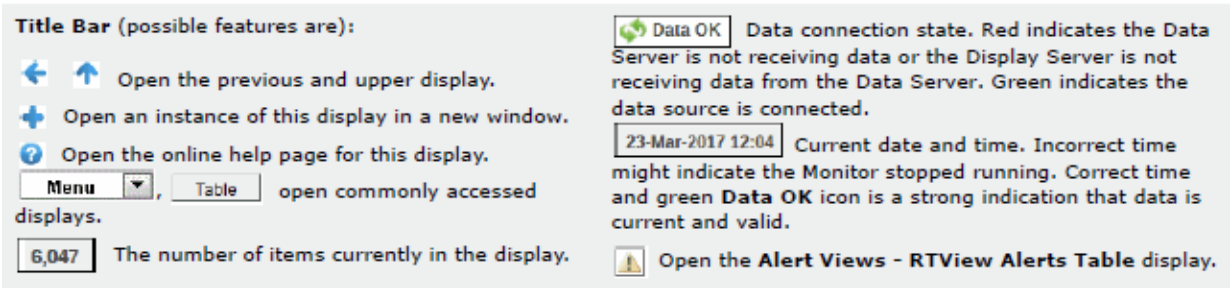

Filter By:

The display might include these filtering options:

- **Cluster:** Select a cluster for which you want to see metrics.
- **Node:** Select a node for which you want to see metrics.
- **Last Update** The date and time the data was last updated.

Node-Level Event Statistics Table:

Each row in the table is a different event, with data in the row columns describing the event.

Note: Row columns in this table with an asterisk (*) at the end of the column definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these columns.

- **Event Count:** The total number of events in the table.
- **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/sec	The rate of "get" operations for the event.
Puts/sec	The rate of "put" operations for the event.
Removes/sec	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.
Timestamp	The date and time, relative to the Data Server, that data was last collected for the engine.

Event Hit Summary

View detailed event performance metrics for a single cluster or **All Clusters,** a node, and an event.

NOTE: Events cause rules to execute in the BusinessEvents Rete network. Events can be created by external phenomena, such as the arrival of a JMS message, or internally when rules are processed. When an event enters the Rete network, it causes a run-to-completion cycle which continues until no further rules can be processed. Each named event that can be handled by a BusinessEvents application is specified at build time in BusinessEvents studio. For details, refer to TIBCO documentation.

Choose a single cluster or **All Clusters**, a node and an event from the drop-down menus.

Title Bar (possible features are):

Open the previous and upper display.

Open an instance of this display in a new window.

Open the online help page for this display.

Menu, Table open commonly accessed displays.

6,047 The number of items currently in the display.

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

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

Open the **Alert Views - RTView Alerts Table** display.

Filter By:

The display might include these filtering options:

Note: Fields in this display with an asterisk (*****) at the end of the field definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these fields.

Cluster:	Select a cluster containing the node and event for which you want to see metrics.
Node:	Select a node containing the event for which you want to see metrics.
Event	Select the event for which you want to see metrics.
Last Update	The date and time in which the data was last updated.
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.
Cache Size	The size of the cache.*****

Handles in Store		The number of handles in the Backing Store for the event.*
Avg Cache Access Time	**Get**	The average time taken for a "get" operation.*
	Put	The average time taken for a "put" operation.*
	Remove	The average time taken for a "remove" operation.*
Cache Access Stats	**Get -- Total Hits**	The total number of "get" operations for the event.*
	Get--Hits	The number of "get" operations for the event since the last data update.*
	Get--Hits/sec	The rate of "get" operations for the event.
	Put--Total Hits	The total number of "put" operations for the event.*
	Put--Hits	The number of "put" operations for the event since the last data update.*
	Put--Hits/sec	The rate of "put" operations for the event.
	Remove--Total Hits	The total number of "remove" operations for the event.*
	Remove--Hits	The number of "remove" operations for the event since the last data update.*
	Remove--Hits/sec	The rate of "remove" operations for the event.
Cache Access Trends		Shows metrics for the selected cluster/node/event combination: **Gets(/sec)** -- Traces the number of "gets" (or rate of "gets" per second depending on **Use Rates** setting) for the event. **Puts(/sec)**-- Traces the number of "puts" (or rate of "puts" per second depending on **Use Rates** setting) for the event. **Removes(/sec)**-- Traces the number of "removes" (or rate of "removes" per second depending on **Use Rates** setting) 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.
	Base at Zero	When this option is checked, zero is set as the Y axis minimum for all graph traces.
	Time Range	Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the button.

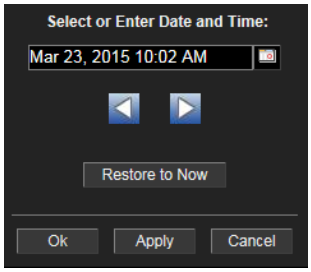

By default, the time range end point is the current time. To change the time range end point, click the button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

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

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

Concept Cache Hits

View a list of concepts and their run-time statistics. Choose a single cluster or **All Clusters** and a node from the drop-down menus.

Sort the table columns when all the rows cannot fit on the screen. For example, sort the **Expired** column so that all expired nodes are listed at the top.

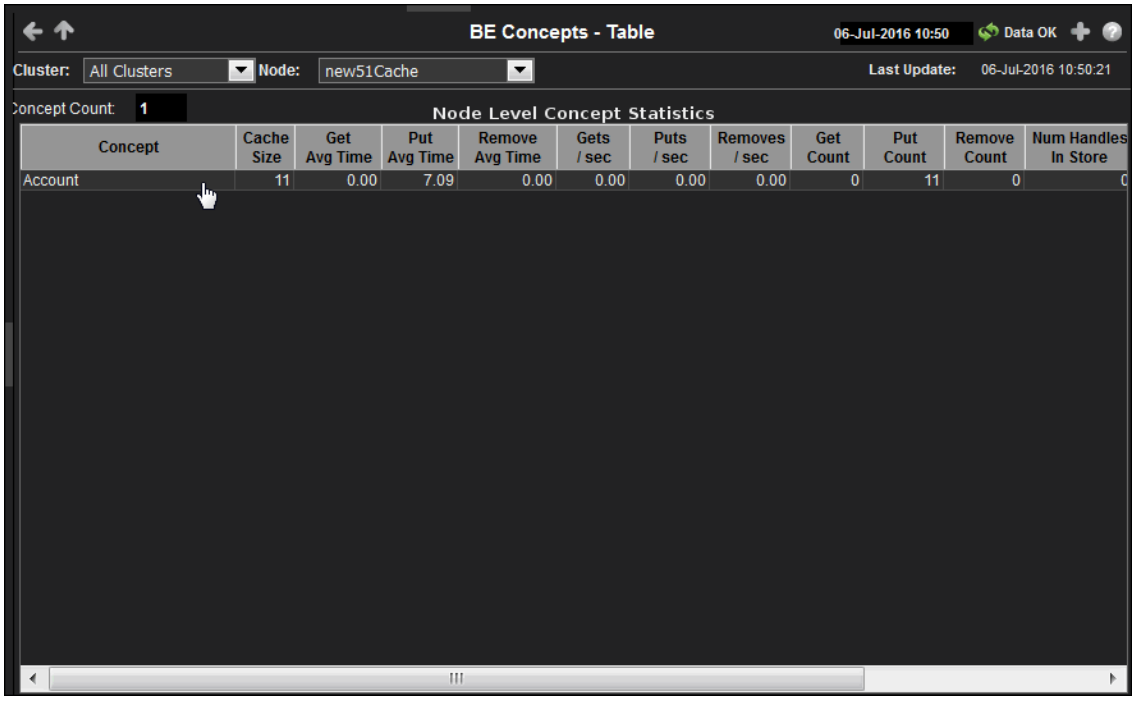

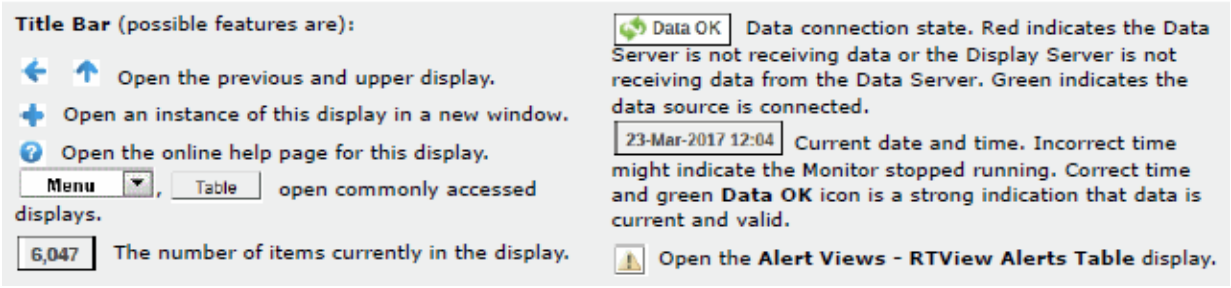

Filter By:

The display might include these filtering options:

Cluster:	Choose a cluster to see metrics for.
Node:	Choose a node to see metrics for.
Last Update	The date and time the data was last updated.

Node-Level Concept Statistics Table:

Each row in the table provides statistics regarding data access for a given BusinessEvents concept.

Note: Row columns in this table with an asterisk (*****) at the end of the column definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these columns.

Concept Count:	The total number of concepts in the table.
Concept	The name of the concept.
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/sec	The rate of "gets" for the concept.
Puts/sec	The rates of "puts" for the concept.
Removes/sec	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.
Timestamp	The date and time, relative to the Data Server, that data was last collected for the concept.

Concept Hit Summary

Use this display to view current and historic data for a single concept. Data in this display can be useful if your BusinessEvents system uses Cache object management. When Cache object management is used, concepts with a sufficiently long time to live (TTL) setting are cached.

Cache reference patterns for certain concepts may be related to incoming events (for example, customer purchase orders with associated inventory queries). The trend charts

show the cache activity of such concepts, and might be useful in diagnosing the behavior of your application over time.

Choose a single cluster or **All Clusters**, a node and a concept from the drop-down menus. Change the trend graph **Time Range** to "zoom in" on the graph and see more detail or "zoom out" from the graph to see larger trends over time.

Title Bar (possible features are):

Open the previous and upper display.

Open an instance of this display in a new window.

Open the online help page for this display.

Menu, Table open commonly accessed displays.

6,047 The number of items currently in the display.

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

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

Open the **Alert Views - RTView Alerts Table** display.

Filter By:

The display might include these filtering options:

Cluster: Select a cluster containing the node and concept for which you want to see metrics.

Node: Select a node containing the concept for which you want to see metrics.

Concept Select the concept for which you want to see metrics.

Fields and Data:

Note: Fields in this table with an asterisk (*) at the end of the field definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these fields.

Last Update The date and time in which the data was last updated in the display.

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.
Cache Size		The size of the cache.*****
Handles in Store		The number of handles in the Backing Store.*****
Avg Cache Access Time	**Get**	The average time taken for a "get" operation.*****
	Put	The average time taken for a "put" operation.*****
	Remove	The average time taken for a "remove" operation.*****
Cache Access Stats	**Get -- Total Hits**	The total number of "get" operations for the concept.*****
	Get--Hits	The number of "get" operations for the concept since the last data update.*****
	Get-- Hits/sec	The rate of "get" operations for the concept.
	Put-- Total Hits	The total number of "put" operations for the concept.*****
	Put--Hits	The number of "put" operations for the concept since the last data update.*****
	Put-- Hits/sec	The rate of "put" operations for the concept.
	Remove-- Total Hits	The total number of "remove" operations for the concept.*****
	Remove-- Hits	The number of "remove" operations for the concept since the last data update.*****
	Remove-- Hits/sec	The rate of "remove" operations for the concept.
Cache Access Trends		Shows metrics for the selected cluster/node/concept combination: **Gets(/sec)** -- Traces the number of "get" operations (or rate of "get" operations depending on **Use Rates** setting) for the concept. **Puts(/sec)**-- Traces the number of "put" operations (or rate of "put" operations depending on **Use Rates** setting) for the concept. **Removes(/sec)**-- Traces the number of "remove" operations (or rate of "remove" operations depending on **Use Rates** setting) 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.
	Base at Zero	When this option is checked, zero is set as the Y axis minimum for all graph traces.
	Time Range	Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click the button.

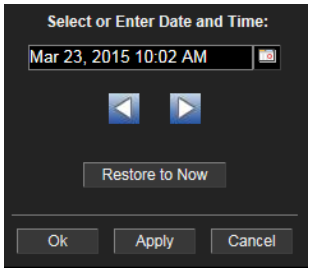

By default, the time range end point is the current time. To change the time range end point, click the button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

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

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

Channels

Use this display to view a list of destinations, which are sources and sinks of events. Destinations are potentially bi-directional, and the table indicates whether events are sent or received.

NOTE: Channels provide a class wrapper for destinations, and make it possible to enable or disable a group of destinations with one operation.

Choose a single cluster or **All Clusters** and a node from the drop-down menus. Each row in the table is a different destination URI. Click a row to view channel details in the **Channels** table.

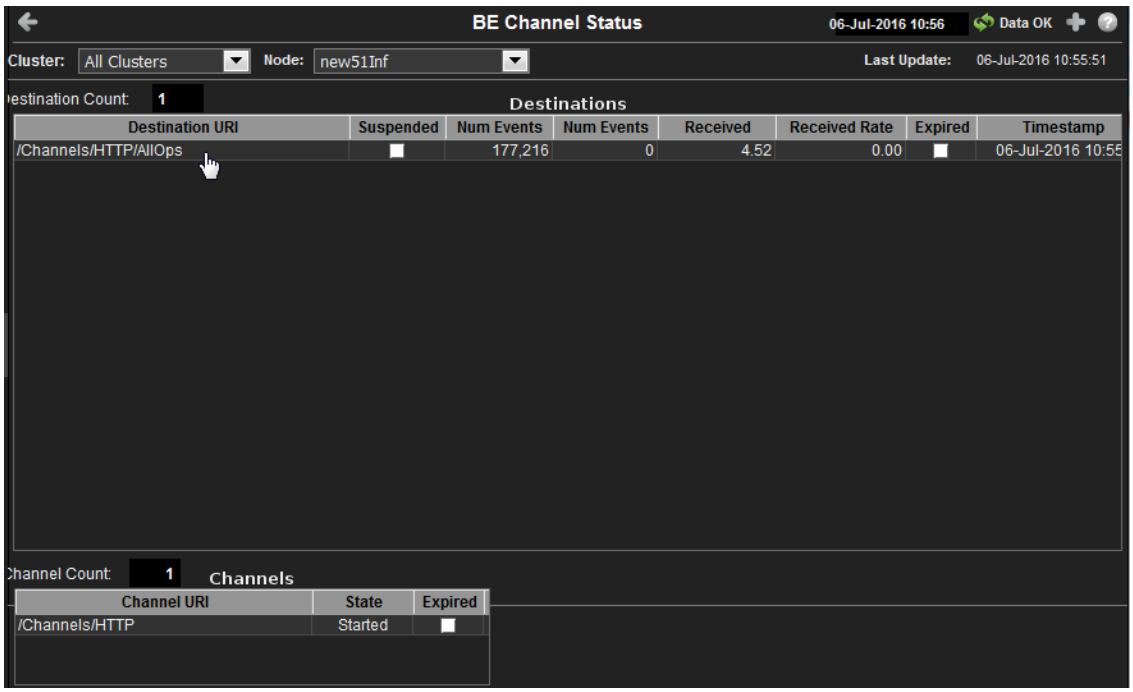

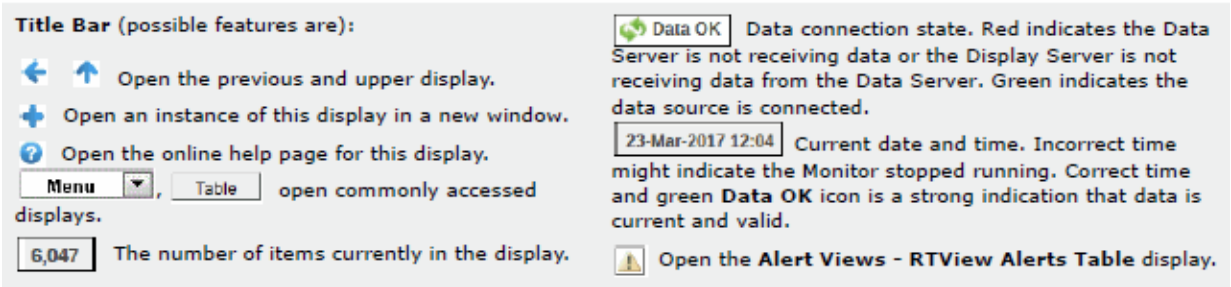

Filter By:

The display might include these filtering options:

Cluster:	Choose a cluster to see metrics for.
Node:	Choose a node to see metrics for.

Destinations Table

Each row in the able provides data for a particular destination.

Note: Row columns in this table with an asterisk (*****) at the end of the column definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these columns.

Destination Count:	The total number of destinations in the table.
Destination URI	The Uniform Resource Identifier (URI) for the destination.*****
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 Rate	The rate of events received by the destination.
Received Rate Last Interval	The rate of events received.
Expired	When checked, performance data has not been received within the time specified (in seconds) in the **Expire Time** field in the **Duration** region in the RTView Configuration Application > (**Project Name**) > **Solution Package Configuration** > **TIBCO BusinessEvents** > **DATA STORAGE** tab. The **Delete Time** field (also in the **Duration** region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
Timestamp	The date and time, relative to the Data Server, that data was last collected for the destination.

Channels Table

Each row in the able 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 Count:	The total number of channels in the table.
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.

All Inference Agents

Use this display to compare agent metrics across deployed engines and verify that the cluster is properly load-balanced. View a list of all the inference agents deployed in each cluster. You can view agent data for a single cluster or all clusters.

The data in this display is identical to the data provided for a single engine in the Cluster Summary display, except that it is aggregated across all inference nodes.

Choose a single cluster or **All Clusters** from the drop-down menus. Each row in the table is a different agent.

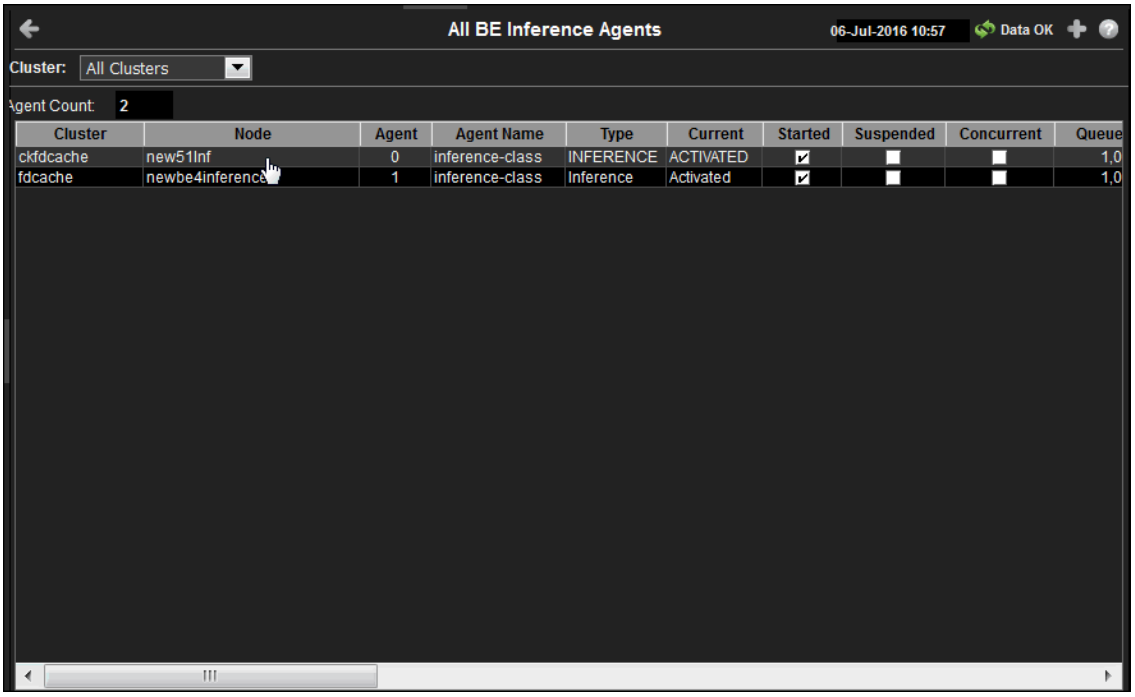

Title Bar (possible features are):

Open the previous and upper display.

Open an instance of this display in a new window.

Open the online help page for this display.

Menu, Table open commonly accessed displays.

6,047 The number of items currently in the display.

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

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

Open the **Alert Views - RTView Alerts Table** display.

Filter By:

The display might include these filtering options:

Cluster: Select the cluster for which you want to see metrics, or select **All Clusters** to see metrics for all clusters.

Table

Each row in the table provides details for an agent.

Note: Row columns in this table with an asterisk (*****) at the end of the column definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these columns.

Agent Count:	The number of agents currently in the table.
Cluster	The name of the TIBCO BusinessEvents cluster.
Node	The name of the node.
Agent ID	A unique string that identifies the agent.
Agent Name	The name of the agent.
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/sec	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.

Timestamp The date and time, relative to the Data Server, that data was last collected for the destination.

All RTC Reports

Use this display to compare RTC metrics across deployed engines. View a list of all the inference engine RTC reports. You can view reports for a single cluster or all clusters.

The data in this display is identical to the data provided for a single engine in the Cluster Summary display, except that it is aggregated across all inference nodes.

Choose a single cluster or **All Clusters** from the drop-down menus. Each row in the table is a different node.

Title Bar (possible features are):

Open the previous and upper display.

Open an instance of this display in a new window.

Open the online help page for this display.

Menu, Table open commonly accessed displays.

6,047 The number of items currently in the display.

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

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

Open the **Alert Views - RTView Alerts Table** display.

Filter By:

The display might include these filtering options:

Cluster: Select the cluster for which you want to see metrics, or select **All Clusters** to see metrics for all clusters.

RTC Txn Manager Reports Table

Each row in the table is a different report. Data in the row columns describe the report.

Note: Row columns in this table with an asterisk (*****) at the end of the column definition contain data that is provided by the TIBCO MBean interface. Refer to TIBCO documentation for more information regarding these columns.

Report Count:	The number of reports currently in the table.
Cluster	The name of the TIBCO BusinessEvents cluster.
Node	The name of the node.
Avg Action Txn Millisec	The average amount of time taken for an action transaction, in milliseconds.*****
Avg Cache Queue Wait Time Millisec	The average cache queue wait time, in milliseconds.*****
Avg Cache Txn Millisec	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 Millisec	The average database queue wait time, in milliseconds.*****
Avg DB Txn Millisec	The average amount of time taken for a database transaction, in milliseconds.*****
Avg Successful Txn Time Millisec	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.
Timestamp	The date and time, relative to the Data Server, that data was last collected for the destination.

Alert Views

These displays present detailed information about all alerts that have occurred in your system. These displays present performance data for your BusinessEvents system. Displays in this View are:

- Alert Detail Table

Alert Detail Table

Use this display to track and manage all alerts that have occurred in the system, add comments, acknowledge or assign Owners to alerts.

Each row in the table is a different active alert. Select one or more rows, right-click and choose **Alert** to see all actions that you can perform on the selected alert(s). Choose **Alert / Set Filter Field** to apply the selected cell data to the **Field Filter** and **Search Text** fields. Or enter filter criteria directly in the **Field Filter** and **Search Text** fields. Click **Clear** to clear the **Field Filter** and **Search Text** fields. Click Sort to order column data.

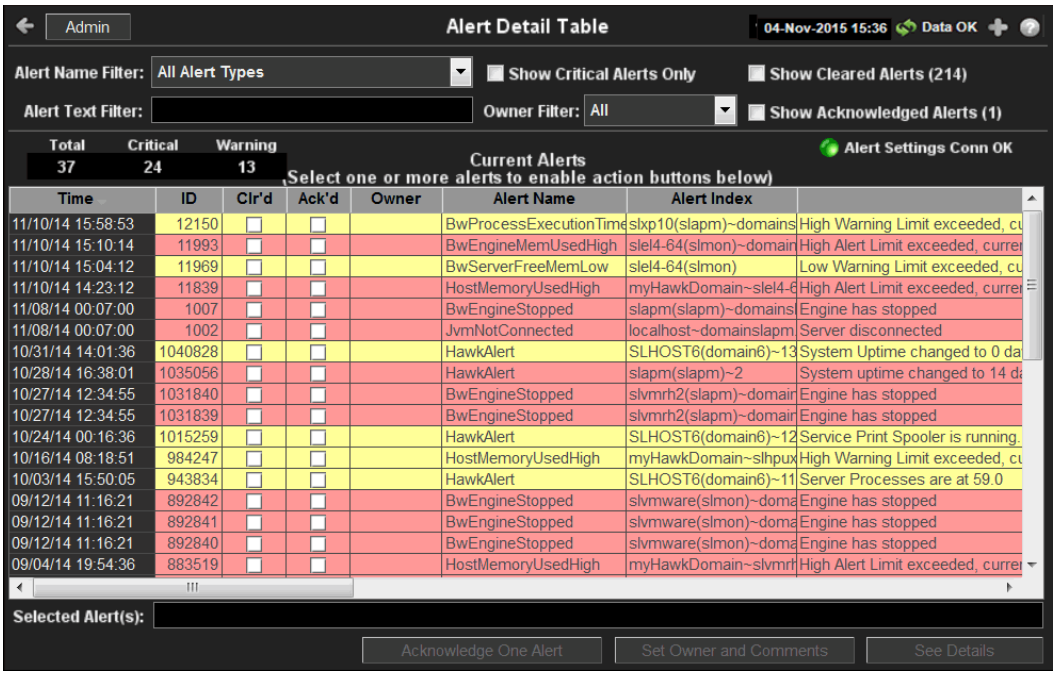

Title Bar (possible features are):

Open the previous and upper display.

Open an instance of this display in a new window.

Open the online help page for this display.

Menu, Table open commonly accessed displays.

6,047 The number of items currently in the display.

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

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

Open the **Alert Views - RTView Alerts Table** display.

Row Color Code:

Tables with colored rows indicate the following:

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

Fields and Data

This display includes:

Alert Name Filter	Select from a list of alert types or select All Alert Types. Filters limit display content and drop down menu selections to only those items that pass through the selected filter's criteria. Therefore if no items match the filter, you may see nothing in a given display and may not have any options available in the drop-down menu(s). **NOTE:** Filter selection is disabled on drill down summary displays.
Show Critical Alerts Only	If selected, only currently critical alerts are shown in the table. Otherwise, all active alerts are shown in the table.
Show Cleared Alerts	If selected, cleared alerts are shown in the table.
Alert Text Filter	Enter all or part of the Alert Text to view specific alerts. For example, High selects and displays all alerts that include High in the Alert Text. **NOTE:** Wild card characters are supported.
Owner Filter	Select the alert **Owner** to show alerts for in the table. **All** Shows alerts for all Owners in the table: **Not Owned** and **Owned By Me** alerts. **Not Owned** Shows only alerts without Owners in the table. **Owned By Me** Shows only alerts for the current user in the table.
Show Acknowledged Alerts	If selected, acknowledged alerts are shown in the table.
Total	Total number of alerts.
Critical	Number of critical alerts.
Warning	Total number of alerts that are currently in a warning state.
Alert Settings Conn OK	The Alert Server connection state: Disconnected. Connected.

Alerts Table

This table lists all active alerts for the current filters.

Time	The time (Java format) that the alert was activated.
ID	A unique string identifier assigned to each activated alert.
Clr'd	When checked, this typically indicates that the alert has been resolved. An alert is automatically cleared when the value being monitored no longer in the alert threshold.
Ack'd	When checked, this typically indicates that the alert is being addressed.
Owner	The named owner assigned by the administrator.
Alert Name	The name of the alert. For a list of all alerts, see Alert Administration.
Alert Index	The IP address and port number for the source (application, server, and so forth) associated with the alert.
Alert Text	Descriptive text about the alert.
Severity	The severity of the alert: **0 = Normal** **1** = Warning / Yellow **2 = Alarm / Red** The color for the alert severity is shown by the row in the alert table.
Source	Name of RTView Data Server sending this data (or localhost).

Selected Alerts Lists the alerts selected in the table.

Acknowledge One Alert	Select one alert from the Current Alerts table and click to acknowledge.
Acknowledge Multiple Alerts	Select one or more alerts from the Current Alerts table and click to acknowledge.
Set Owner and Comments	Select one or more alerts from the Current Alerts table and click to open the Set Owner and Comments dialog.
See Details	Select an alert from the Current Alerts table and click to open the Set Owner and Comments dialog.

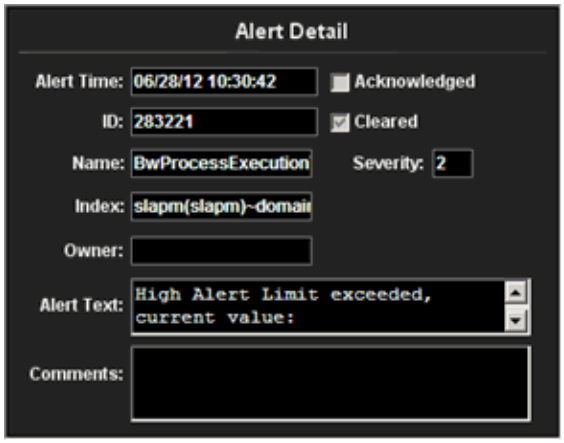

Administration

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

- Alert Administration
- Tabular Alert Administration
- Setting Override Alerts
- Alert Administration Audit
- Metrics Administration
- RTView Cache Tables
- Agent Administration
- Version Info
- About

Alert Administration

Set global or override alert thresholds. Alert settings are global by default.

The table describes the global settings for all alerts on the system. To filter the alerts listed in the table, enter a string in the **Alert Filter** field and press **<enter>** or click elsewhere in the display. Filters are case sensitive and no wildcard characters are needed for partial strings. For example, if you enter Server in the **Alert Filter** field, it filters the table to show only alerts with **Server** in the name. Choose **Clear** to clear the filter.

Global Thresholds

To set a global alert, select an alert from the **Active Alert Table**. The name of the selected alert populates the **Settings for Selected Alert Name** field. Edit the **Settings for Selected Alert** and click **Save Settings** when finished.

The manner in which global alerts are applied depends on the Solution Package. For example, the EMS Monitor Solution Package has queue alerts, topic alerts and server alerts. When a

queue alert is applied globally, it is applied to all queues on all servers. Likewise, a server alert applies to all servers, and a topic alert applies to all topics on all servers.

Override Thresholds

Setting override alerts allows you to set thresholds for a single resource (for example, a single server). Override alerts are useful if the majority of your alerts require the same threshold setting, but there are other alerts that require a different threshold setting. For example, you might not usually be concerned with execution time at a process level, but perhaps certain processes are critical. In this case, you can apply alert thresholds to each process individually.

To apply an individual alert you Index the Monitored Instance or resource. The Index Types available are determined by the Solution Package installed. For example, the EMS Monitor package lets you set an alert for a specific *topic* on a specific *server* (such as the PerServerTopic Index option), rather than for all topics on all servers.

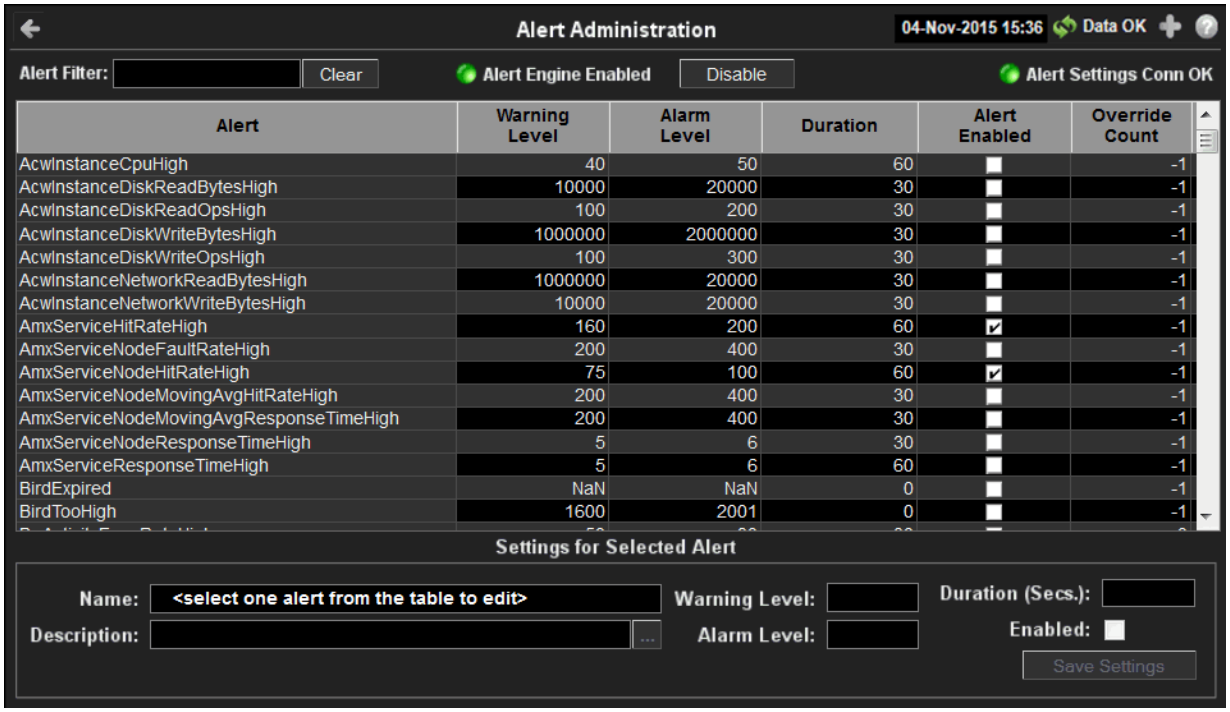

Title Bar (possible features are):

Open the previous and upper display.

Open an instance of this display in a new window.

Open the online help page for this display.

Menu, Table open commonly accessed displays.

6,047 The number of items currently in the display.

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

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

Open the **Alert Views - RTView Alerts Table** display.

Fields and Data
This display includes:

Alert Filter	Enter the (case-sensitive) string to filter the table by the **Alert** table column value. **NOTE:** Partial strings can be used without wildcard characters. Press **<enter>** or click elsewhere in the display to apply the filter.
Clear	Clears the **Alert Filter** entry.
Alert Engine Enabled	Alerting is disabled. Alerting is enabled (by default).
Disable	Suspends all alerting.
Alert Settings Conn OK	The Alert Server connection state: Disconnected. Connected.

Active Alert Table

This table describes the global settings for all alerts on the system. Select an alert. The name of the selected alert populates the **Settings for Selected Alert Name** field (in the lower panel). Edit **Settings for Selected Alert** fields and click **Save Settings**.

NOTE: To filter the alerts shown in the table by Solution Package, use the **$rtvAlertPackageMask** substitution.

Alert	The name of the alert.
Warning Level	The global warning threshold for the selected alert. When the specified value is exceeded a warning is executed.
Alarm Level	The global alarm threshold for the selected alert. When the specified value is exceeded an alarm is executed.
Duration (Secs)	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. **0** is for immediate execution.
Alert Enabled	When checked, the alert is enabled globally.
Override Count	The number of times thresholds for this alert have been defined individually in the **Tabular Alert Administration** display.

Settings for Selected Alert

To view or edit global settings, select an alert from the **Active Alert Table**. Edit the **Settings for Selected Alert** fields and click **Save Settings** when finished.

To set override alerts, click on **Override Settings** to open the **Tabular Alert Administration** display.

Name	The name of the alert selected in the **Active Alert Table**.
Description	Description of the selected alert. Click Calendar ... for more detail.
Warning Level	Set the Global warning threshold for the selected alert. When the specified value is exceeded a warning is executed. To set the warning to occur sooner, reduce the Warning Level value. To set the warning to occur later, increase the Warning Level value. **NOTE:** For low value-based alerts (such as **EmsQueuesConsumerCountLow**), to set the warning to occur sooner, increase the Warning Level value. To set the warning to occur later, reduce the Warning Level value.
Alarm Level	Set the Global alarm threshold for the selected alert. When the specified value is exceeded an alarm is executed. To set the alarm to occur sooner, reduce the Alarm Level value. To set the warning to occur later, increase the Alarm Level value. **NOTE:** For low value-based alerts (such as **EmsQueuesConsumerCountLow**), to set the alarm to occur sooner, increase the Alarm Level value. To set the

	alarm to occur later, reduce the Alarm Level value.
Duration	Set the amount of time (in seconds) that the value must be above the specified Warning Level or Alarm Level threshold before an alert is executed. **0** is for immediate execution. This setting is global.
Enabled	Check to enable alert globally.
Save Settings	Click to apply alert settings.
Override Settings	Click to open the **Tabular Alert Administration** display to set override alerts on the selected alert.

Tabular Alert Administration

Set override alerts (override global alert settings). This display opens when you select an alert in the **Alert Administration** display and then select **Override Settings**.

For step-by-step instructions setting thresholds for individual alerts, see **Setting Override Alerts**.

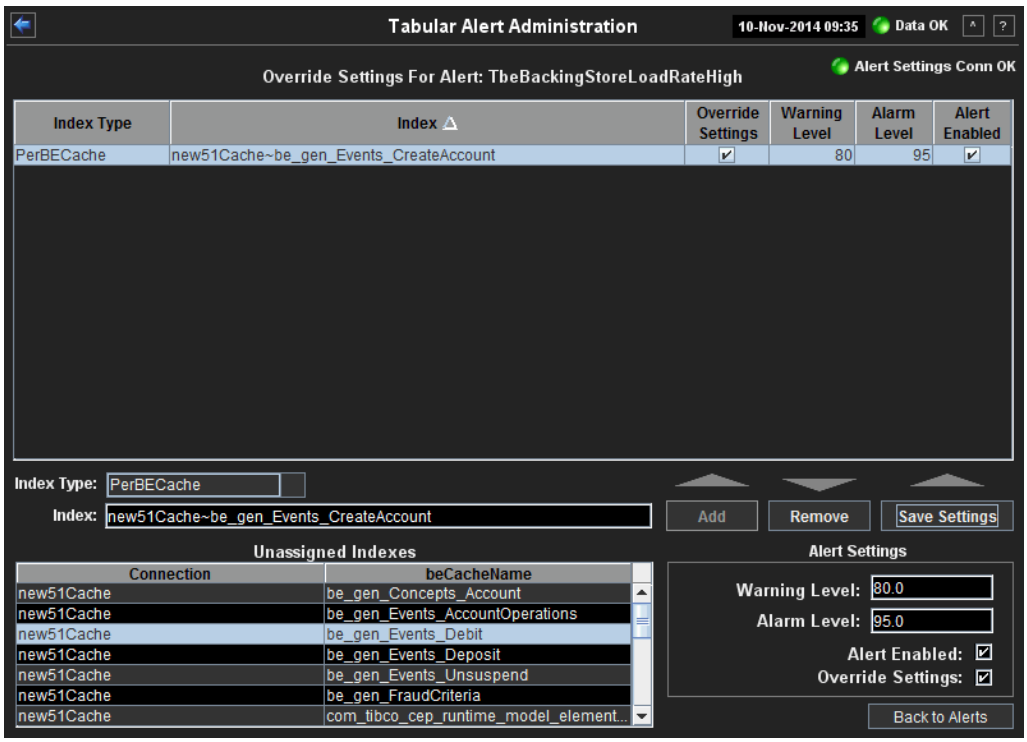

Fields and Data
This display includes:

Alert Settings Conn OK
The connection state.
- No servers are found.
- One or more servers are delivering data.

Override Settings For Alert:(name)

This table lists and describes alerts that have override settings for the selected alert. Select a row to edit alert thresholds. The selected item appears in the **Index** field. Edit settings in the **Alert Settings** fields, then click **Save** Settings.

Index Type	Select the type of alert index to show in the **Values** table. Options in this drop-down menu are populated by the type of alert selected, which are determined by the Package installed. For example, with the EMS

		Monitor package the following Index Types are available: • PerServer: Alert settings are applied to a specific server. • PerQueue: Alert settings are applied to the queue on each server that has the queue defined. • PerServerQueue: Alert settings are applied to a single queue on a specific server. • PerTopic: Alert settings are applied to the topic on each server that has the topic defined. • PerServerTopic: Alert settings are applied to a single topic on a specific server.
	Index	The value of the index column.
	Override Settings	When checked, the override settings are applied.
	Alert Enabled	When checked, the alert is enabled.
Index Type		Select the index type. The index type specifies how to apply alert settings. For example, to a queue (topic or JVM, and so forth) across all servers, or to a queue on a single server. **NOTE:** Options in this drop-down menu are populated by the type of alert selected from the **Alert Administration** display. Index Types available depend on the Package installed.
Index		The selected index column to be edited. This field is populated by the selection made in the **Unassigned Indexes** table.
Unassigned Indexes		This table lists all possible indexes corresponding to the Index Type chosen in the drop-down list. Select a row to apply individual alert thresholds. The selected item appears in the **Index** field. Edit settings in the **Alert Settings** fields, then click **Add**.
Add		Click to add changes made in **Alert Settings**, then click **OK** to confirm.
Remove		Click to remove an alert selected in the **Index Alert Settings** table, then click **OK** to confirm.
Save Settings		Click to save changes made to alert settings.

Alert Settings

Select a topic, server or queue from the **Unassigned Indexes** table and edit the following settings.

	Warning Level	Set the warning threshold for the selected alert. When the specified value is exceeded a warning is executed. To set the warning to occur sooner, reduce the Warning Level value. To set the warning to occur later, increase the Warning Level value. **NOTE:** For low value-based alerts (such as **EmsQueuesConsumerCountLow**), to set the warning to occur sooner, increase the Warning Level value. To set the warning to occur later, reduce the Warning Level value. Click **Save Settings** to save settings.
	Alarm Level	Set the alarm threshold for the selected alert. When the specified value is exceeded an alarm is executed. To set the alarm to occur sooner, reduce the Alarm Level value. To set the warning to occur later, increase the Alarm Level value. NOTE: For low value-based alerts (such as **EmsQueuesConsumerCountLow**), to set the alarm to occur sooner, increase the Alarm Level value. To set the alarm to occur later, reduce the Alarm Level value. Click **Save Settings** to save settings.
	Alert Enabled	Check to enable the alert, then click **Save Settings**.
	Override Settings	Check to enable override global setting, then click **Save Settings**.
Back to Alerts		Returns to the **Administration** - **Alert Administration** display.

Setting Override Alerts

Perform the following steps to set an override alert. Index Types available depend on the Solution Package installed. In this example, we use the EMS Monitor Package to illustrate.

NOTE: To turn on an alert, both Alert Enabled and Levels Enabled must be selected.

To turn on/off, change threshold settings, enable/disable or remove an alert on a single resource:

1. In the **Alert Administration** display, select an alert in the **Active Alert Table** and click **Edit Index Levels**. The **Tabular Alert Administration** display opens.

2. In the **Tabular Alert Administration** display, from the **Index Type** drop-down menu, select the Index type (options are populated by the type of alert you previously selected). For example, with the EMS Monitor package, select PerServerQueue, PerServerTopic or PerServer. **NOTE:** If you select PerServerQueue or PerServerTopic, the alert settings are applied to the queue or topic on a single server.

3. In the **Values** table, select the server to apply alert settings and click **Add**. In a few moments the server appears in the **Index Alert Settings** table.

4. In the **Index Alert Settings** table select the server.

5. In the **Alert Settings** panel (lower right), if needed, modify the **Warning Level** and **Alarm Level** settings.

6. In the **Alert Settings** panel, set the following as appropriate.

   To turn on the alert for this index with the given thresholds:

   Alert Enabled Select this option.

   Levels Enabled Select this option.

   To turn off the alert for only this index (global alert thresholds will no longer apply to this index):

   Alert Enabled Deselect this option.

   Levels Enabled Select this option.

   To no longer evaluate this indexed alert and revert to global settings (or, optionally, Remove it if it is never to be used again):

   Alert Enabled Not used.

   Levels Enabled Deselect this option.

7. Click **Save Settings**. In a few moments the modifications are updated in the **Index Alert Settings** table.

Alert Administration Audit

View alert management details such as alert threshold modifications.

Each table row is a single modification made to an alert. To view modifications for a single alert in a group, sort the **ALERTNAME** column using the button.

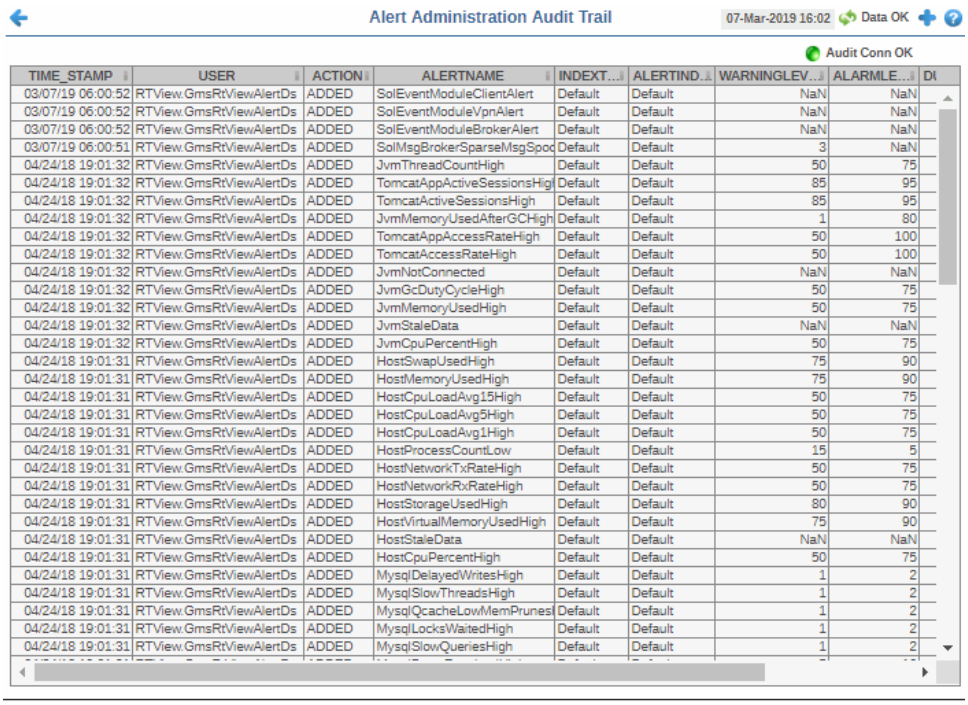

TIME_STAMP	USER	ACTION	ALERTNAME	INDEXT...	ALERTIND...	WARNINGLEV...	ALARMLE...	DI
03/07/19 06:00:52	RTView.GmsRtViewAlertDs	ADDED	SolEventModuleClientAlert	Default	Default	NaN	NaN	
03/07/19 06:00:52	RTView.GmsRtViewAlertDs	ADDED	SolEventModuleVpnAlert	Default	Default	NaN	NaN	
03/07/19 06:00:52	RTView.GmsRtViewAlertDs	ADDED	SolEventModuleBrokerAlert	Default	Default	NaN	NaN	
03/07/19 06:00:51	RTView.GmsRtViewAlertDs	ADDED	SolMsgBrokerSparseMsgSpoo	Default	Default	3	NaN	
04/24/18 19:01:32	RTView.GmsRtViewAlertDs	ADDED	JvmThreadCountHigh	Default	Default	50	75	
04/24/18 19:01:32	RTView.GmsRtViewAlertDs	ADDED	TomcatAppActiveSessionsHig	Default	Default	85	95	
04/24/18 19:01:32	RTView.GmsRtViewAlertDs	ADDED	TomcatActiveSessionsHigh	Default	Default	85	95	
04/24/18 19:01:32	RTView.GmsRtViewAlertDs	ADDED	JvmMemoryUsedAfterGCHigh	Default	Default	1	80	
04/24/18 19:01:32	RTView.GmsRtViewAlertDs	ADDED	TomcatAppAccessRateHigh	Default	Default	50	100	
04/24/18 19:01:32	RTView.GmsRtViewAlertDs	ADDED	TomcatAccessRateHigh	Default	Default	50	100	
04/24/18 19:01:32	RTView.GmsRtViewAlertDs	ADDED	JvmNotConnected	Default	Default	NaN	NaN	
04/24/18 19:01:32	RTView.GmsRtViewAlertDs	ADDED	JvmGcDutyCycleHigh	Default	Default	50	75	
04/24/18 19:01:32	RTView.GmsRtViewAlertDs	ADDED	JvmMemoryUsedHigh	Default	Default	50	75	
04/24/18 19:01:32	RTView.GmsRtViewAlertDs	ADDED	JvmStaleData	Default	Default	NaN	NaN	
04/24/18 19:01:32	RTView.GmsRtViewAlertDs	ADDED	JvmCpuPercentHigh	Default	Default	50	75	
04/24/18 19:01:31	RTView.GmsRtViewAlertDs	ADDED	HostSwapUsedHigh	Default	Default	75	90	
04/24/18 19:01:31	RTView.GmsRtViewAlertDs	ADDED	HostMemoryUsedHigh	Default	Default	75	90	
04/24/18 19:01:31	RTView.GmsRtViewAlertDs	ADDED	HostCpuLoadAvg15High	Default	Default	50	75	
04/24/18 19:01:31	RTView.GmsRtViewAlertDs	ADDED	HostCpuLoadAvg5High	Default	Default	50	75	
04/24/18 19:01:31	RTView.GmsRtViewAlertDs	ADDED	HostCpuLoadAvg1High	Default	Default	50	75	
04/24/18 19:01:31	RTView.GmsRtViewAlertDs	ADDED	HostProcessCountLow	Default	Default	15	5	
04/24/18 19:01:31	RTView.GmsRtViewAlertDs	ADDED	HostNetworkTxRateHigh	Default	Default	50	75	
04/24/18 19:01:31	RTView.GmsRtViewAlertDs	ADDED	HostNetworkRxRateHigh	Default	Default	50	75	
04/24/18 19:01:31	RTView.GmsRtViewAlertDs	ADDED	HostStorageUsedHigh	Default	Default	80	90	
04/24/18 19:01:31	RTView.GmsRtViewAlertDs	ADDED	HostVirtualMemoryUsedHigh	Default	Default	75	90	
04/24/18 19:01:31	RTView.GmsRtViewAlertDs	ADDED	HostStaleData	Default	Default	NaN	NaN	
04/24/18 19:01:31	RTView.GmsRtViewAlertDs	ADDED	HostCpuPercentHigh	Default	Default	50	75	
04/24/18 19:01:31	RTView.GmsRtViewAlertDs	ADDED	MysqlDelayedWritesHigh	Default	Default	1	2	
04/24/18 19:01:31	RTView.GmsRtViewAlertDs	ADDED	MysqlSlowThreadsHigh	Default	Default	1	2	
04/24/18 19:01:31	RTView.GmsRtViewAlertDs	ADDED	MysqlQcacheLowMemPrunes	Default	Default	1	2	
04/24/18 19:01:31	RTView.GmsRtViewAlertDs	ADDED	MysqlLocksWaitedHigh	Default	Default	1	2	
04/24/18 19:01:31	RTView.GmsRtViewAlertDs	ADDED	MysqlSlowQueriesHigh	Default	Default	1	2	

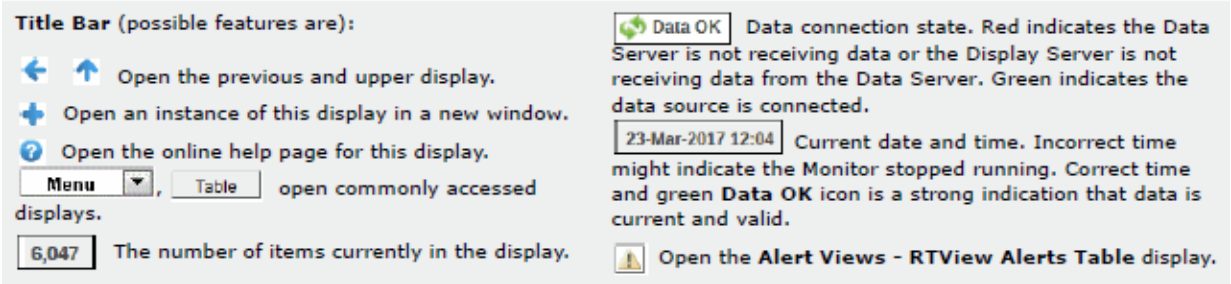

Audit Conn OK	The Alert Server connection state: Disconnected. Connected.
TIME_STAMP	The date and time of the modification.
USER	The user name of the administrator who made the modification.
ACTION	The type of modification made to the alert, such as UPDATED.
ALERTNAME	The name of the alert modified.
INDEXTYPE	The type of alert Index.
ALERTINDEX	The IP address and port number for the source (application, server, and so forth) associated with the alert.
WARNINGLEVEL	The warning threshold value for the alert at the time this modification was made, as indicated in the **TIME_STAMP** column. The warning level is a threshold that, when exceeded, a warning is executed.
ALARMLEVEL	The alarm threshold value for the alert at the time this modification was made, as indicated in the **TIME_STAMP** column. The alarm level is a threshold that, when exceeded, an alarm is executed.
DURATION	The duration value for the alert at the time this modification was made, as indicated in the **TIME_STAMP** column. The alert duration is the amount of time (in seconds) that a value must exceed the specified Warning Level or Alarm Level threshold before an alert is executed. 0 is for immediate execution.

ENABLED	When checked, indicates the alert was Enabled at the time this modification was made, as indicated in the **TIME_STAMP** column.
USEINDEX	When checked, this action was performed on an override alert (the alert does not use the global settings).

Metrics Administration

Verify when TIBCO metrics were last queried by the Monitor. The data in this display is predominantly used for debugging by Technical Support.

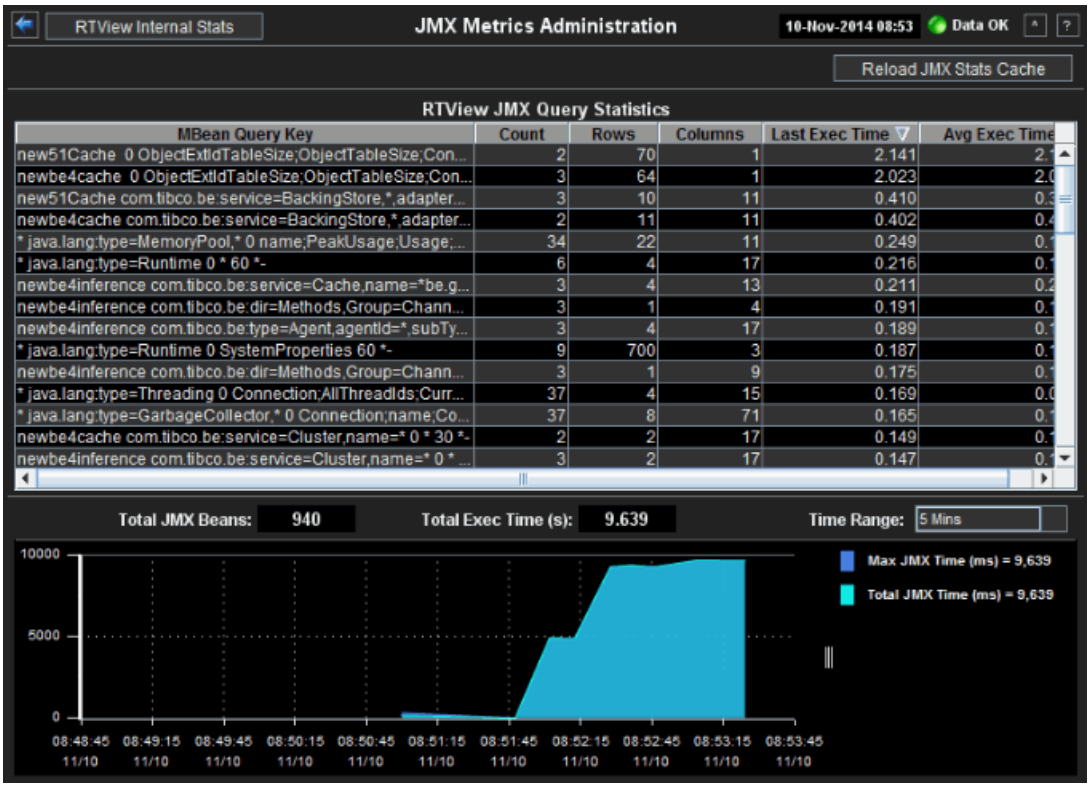

Title Bar: Indicators and functionality might include the following:

Open the previous and upper display.

Table Navigate to displays commonly accessed from this display.

23-Mar-2017 12:04 The current date and time. When the time is incorrect, this might indicate that RTView stopped running. When the time is correct and the **Data OK** indicator is green, this is a strong indication that the platform is receiving current and valid data.

Data OK The data connection state. Red indicates the data source is disconnected (for example, the Data Server is not receiving data, or the Display Server is not receiving data from the Data Server). Green indicates the data source is connected.

Open the **Alert Views - RTView Alerts Table** display.

Open an instance of this display in a new window.

Open the online help page for this display.

Reload JMX Stats Cache	Reloads the cache.
MBean Query Key	The string used for the data attachment to this data object.

Count	The number of listeners for this data object. For example, graphical objects and function arguments.
Rows	The number of rows (of MBean data that match the query) returned from JMX Query using MBean Query Key.
Columns	The number of columns (of MBean data that match the query) returned from JMX Query using MBean Query Key.
Last Exec Time	The last time a query was executed for the metric associated with this data object.
Avg Exec Time	The average amount of time for a query to execute for the metric associated with this data object.
Total JMX Beans	The total number of MBeans returned by all active JMX Queries (this should be the sum of the **Rows** column).
Total Exec Time	The total amount of time to execute all the JMX Queries (this should be approximately the sum of the **Last Exec Time** column).
Trend Graph	-- Traces the following: **Max JMX Time** -- Traces the maximum execution time, in seconds, for all MBean Query Keys in the table. **Total JMX Time** -- Traces the cumulative execution time, in seconds, for all MBean Query Keys in the table.

Time Range

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

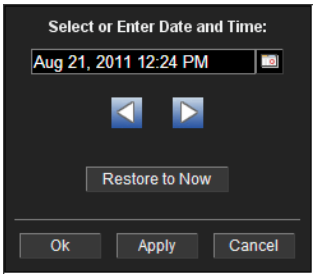

To change the time range select Open Calendar , choose the date and time, then click **OK**. Or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM:ss**. For example, Aug 21, 2011 12:24 PM. Click **Apply**. Use the Navigation Arrows to move forward or backward one 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.

RTView Cache Tables

View data that RTView is capturing and maintaining. Drill down and view details of RTView Cache Tables. Use this data for debugging. This display is typically used for troubleshooting with Technical Support.

Choose a cache table from the upper table to see cached data.

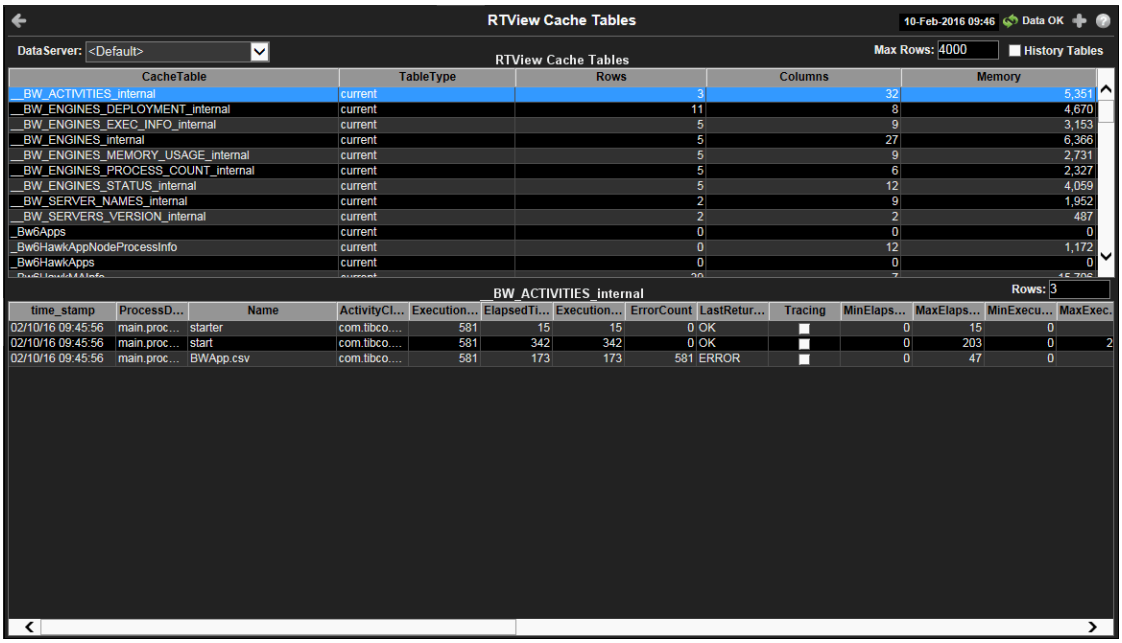

Title Bar (possible features are):

Open the previous and upper display.

Open an instance of this display in a new window.

Open the online help page for this display.

Menu, Table open commonly accessed displays.

6,047 The number of items currently in the display.

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

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

Open the **Alert Views - RTView Alerts Table** display.

DataServer Select a data server from the drop down menu.

Max Rows Enter the maximum number of rows to display in RTView Cache Tables.

History Tables Select to include all defined history tables in RTView Cache Tables.

RTView Cache Tables

This table lists and describes all defined RTView Cache Tables for your system. Cache tables gather Monitor data and are the source that populate the Monitor displays.

NOTE: When you click on a row in RTView Cache Tables a supplemental table will appear that gives more detail on the selected Cache Table.

CacheTable The name of the cache table.

TableType The type of cache table:

- **current** Current table which shows the current values for each index.
- **current_condensed** Current table with primary compaction configured.
- **history** History table.
- **history_condensed** History table with primary compaction configured.

Rows Number of rows currently in the table.

Columns	Number of columns currently in the table.
Memory	Amount of space, in bytes, used by the table.

Agent Administration

Verify when agent metrics were last queried by the Monitor. The data in this display is predominantly used for debugging by Technical Support.

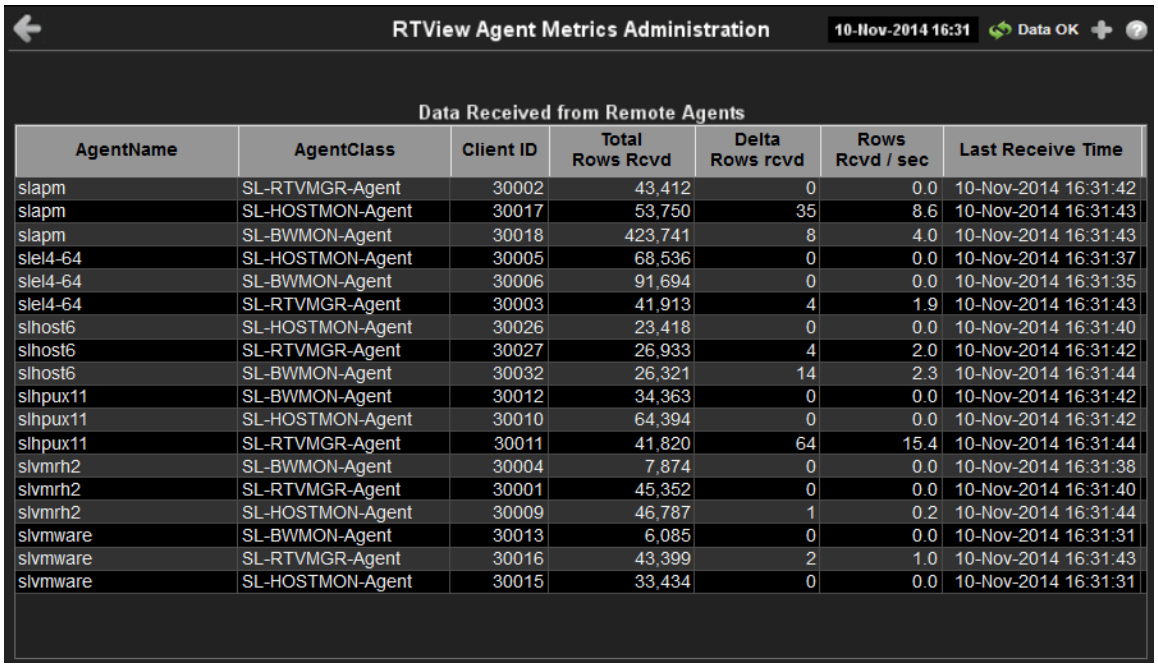

AgentName	AgentClass	Client ID	Total Rows Rcvd	Delta Rows rcvd	Rows Rcvd / sec	Last Receive Time
slapm	SL-RTVMGR-Agent	30002	43,412	0	0.0	10-Nov-2014 16:31:42
slapm	SL-HOSTMON-Agent	30017	53,750	35	8.6	10-Nov-2014 16:31:43
slapm	SL-BWMON-Agent	30018	423,741	8	4.0	10-Nov-2014 16:31:43
slel4-64	SL-HOSTMON-Agent	30005	68,536	0	0.0	10-Nov-2014 16:31:37
slel4-64	SL-BWMON-Agent	30006	91,694	0	0.0	10-Nov-2014 16:31:35
slel4-64	SL-RTVMGR-Agent	30003	41,913	4	1.9	10-Nov-2014 16:31:43
slhost6	SL-HOSTMON-Agent	30026	23,418	0	0.0	10-Nov-2014 16:31:40
slhost6	SL-RTVMGR-Agent	30027	26,933	4	2.0	10-Nov-2014 16:31:42
slhost6	SL-BWMON-Agent	30032	26,321	14	2.3	10-Nov-2014 16:31:44
slhpux11	SL-BWMON-Agent	30012	34,363	0	0.0	10-Nov-2014 16:31:42
slhpux11	SL-HOSTMON-Agent	30010	64,394	0	0.0	10-Nov-2014 16:31:42
slhpux11	SL-RTVMGR-Agent	30011	41,820	64	15.4	10-Nov-2014 16:31:44
slvmrh2	SL-BWMON-Agent	30004	7,874	0	0.0	10-Nov-2014 16:31:38
slvmrh2	SL-RTVMGR-Agent	30001	45,352	0	0.0	10-Nov-2014 16:31:40
slvmrh2	SL-HOSTMON-Agent	30009	46,787	1	0.2	10-Nov-2014 16:31:44
slvmware	SL-BWMON-Agent	30013	6,085	0	0.0	10-Nov-2014 16:31:31
slvmware	SL-RTVMGR-Agent	30016	43,399	2	1.0	10-Nov-2014 16:31:43
slvmware	SL-HOSTMON-Agent	30015	33,434	0	0.0	10-Nov-2014 16:31:31

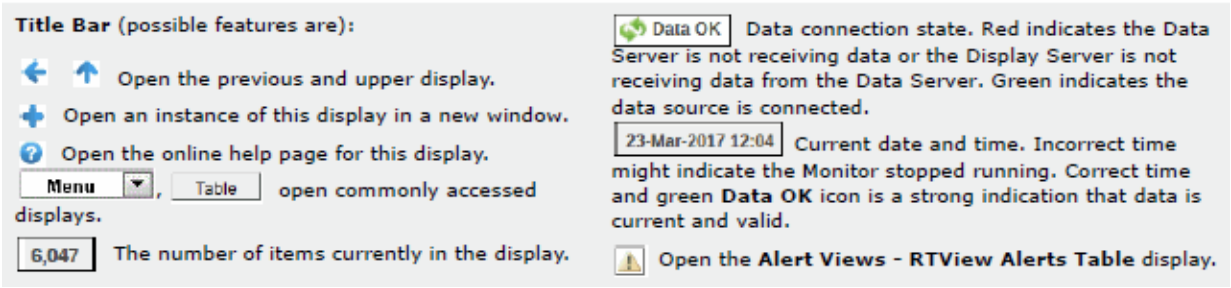

Data Received from Remote Agents Table	
AgentName	Name of the agent.
AgentClass	Class of the agent.
Client ID	Unique client identifier.
Total Rows Rcvd	Total number of rows of data received.
Rows Rcvd/sec	Number of rows of data received per second.
Last Receive Time	Last time data was received from the agent.

Version Info

This display provides detailed version information for all of the connected RTView applications. You can view specific applications by filtering data using the **Source**, **Connection**, **Filter Field**, and **Filter Value** fields at the top of the display. This display provides valuable information about the version of each jar that is used in each connected RTView application that can be used to help Technical Support when issues arise. Rows in the table where the **JarConfiguration** does not match the **ApplicationConfiguration** are highlighted in teal.

> **Note:** RTView applications running versions previous to this enhancement will only have one row in the table and will display "version info not supported in this version" in the **ApplicationConfiguration** column.

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

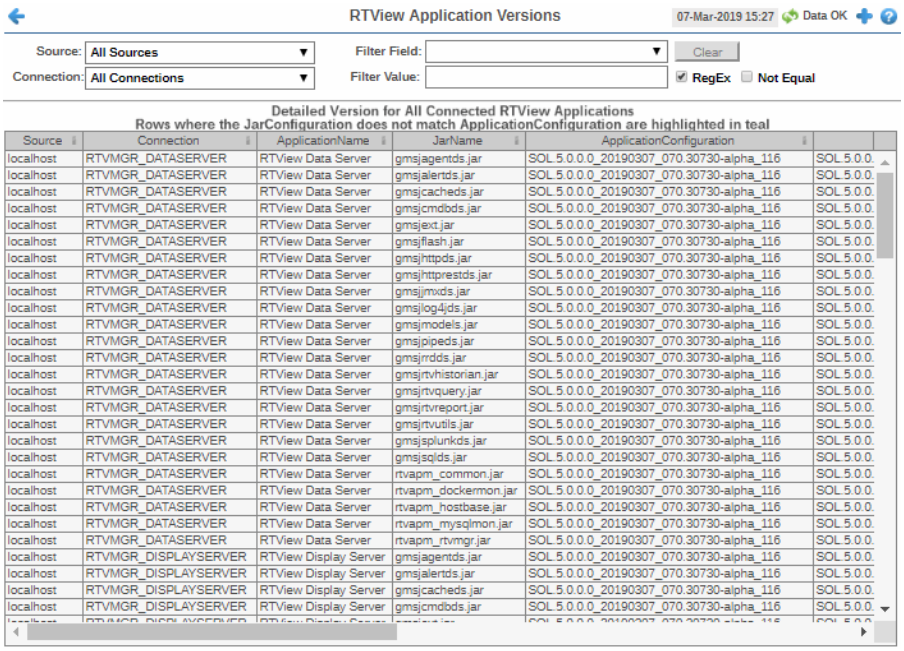

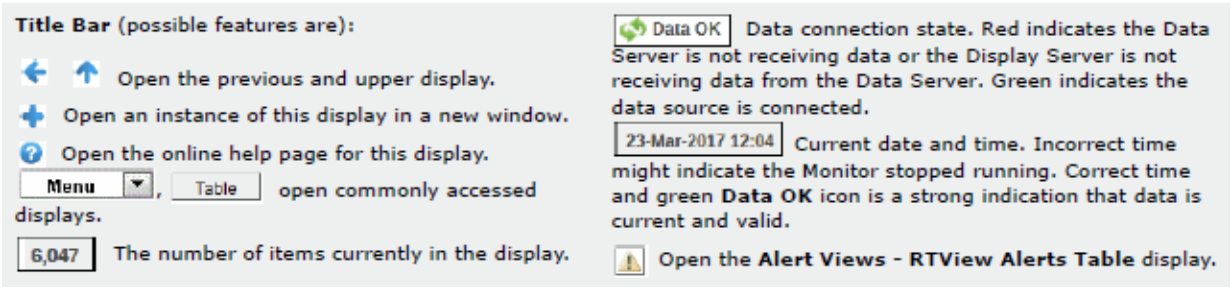

Fields and Data
This display includes:

Source	Select a filter value for the **Source** column.
Connection	Select a filter value for the **Connection** column.
Filter Field	Select a table column from the drop-down menu to perform a search in: **ApplicationName, JarName, ApplicationConfiguration,**

	JarConfiguration, JarVersionNumber, JarVersionDate, JarReleaseDate, and **JarMicroVersion**.
	Filters limit display content and drop-down menu selections to only those items that pass through the selected filter's criteria. If no items match the filter, you might have zero search results (an empty table). Double-clicking on a specific field in the table will populate this field with the selected field's content. For example, double-clicking on the **DataServerName** field in one of the rows displays the entire field's content into this field.
Clear	Clears entries in the **Filter Field** display list, **Filter Value** field, and **Not Equal** check box.
Filter Value	Enter the (case-sensitive) string to search for in the selected **Filter Field**.
RegEx	Select this check box to use the **Filter Value** as a regular expression when filtering. When selected, the **Not Equal** check box displays.
Not Equal	Works in conjunction with the **RegEx** field. Selecting this check box searches for values in the specified **Filter Field** that are NOT equal to the value defined in the **Filter Value** field. For example, if the **Filter Field** specified is **JarMicroVersion**, the **Filter Value** is specified as **317,** and this check box is selected, then only those rows containing **JarMicroVersion** fields NOT EQUAL to **317** will display. This field is only enabled when the **RegEx** check box is checked.
Source	The name of the source of the RTVMGR.
Connection	Lists the name of the JMX connection to the RTView application.
Application Name	Lists the name of the application.
JarName	Lists the name of the jar used in the connected application.
Application Configuration	Lists the configuration string of the application. This string contains the main application version that corresponds to the version information printed to the console at startup.
JarConfiguration	Lists the configuration string for the jar.
JarVersionNumber	Lists the version number for the jar.
JarVersionDate	Lists the version date for the jar.
JarReleaseType	Lists the release type for the jar.
JarMicroVersion	Lists the micro version for the jar.
Expired	When checked, this connection is expired due to inactivity.
time_stamp	The time at which the information in the current row was last received.
DataServerName	The name of the RTVMGR data server connection.

About

This display shows details about the Monitor version and data sources available to your system.

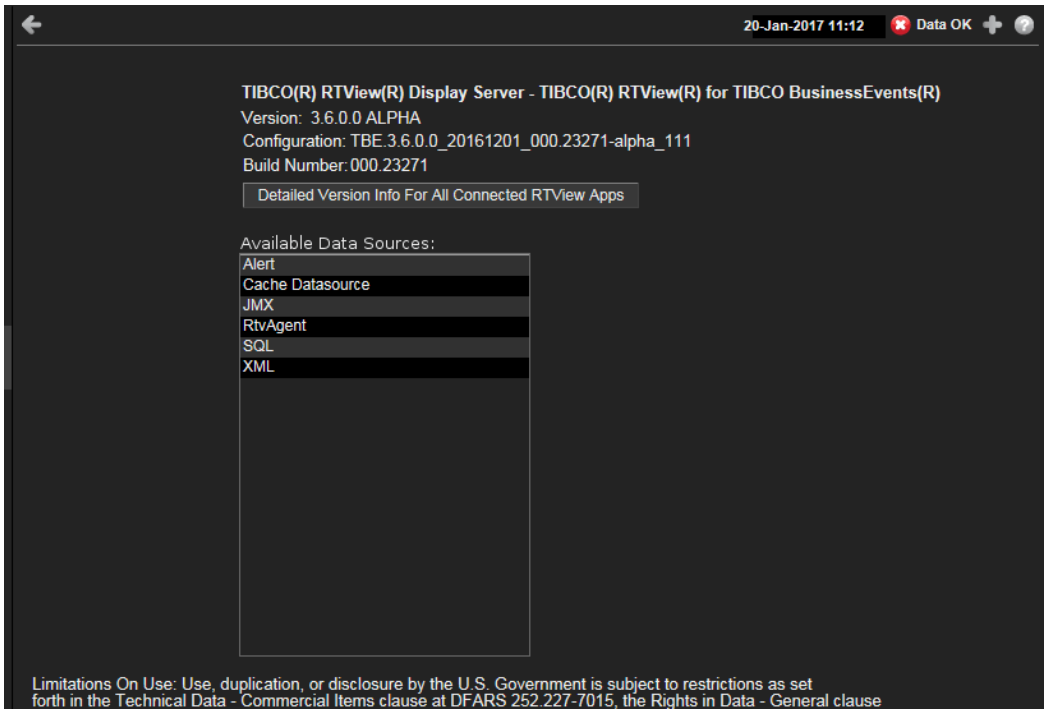
20-Jan-2017 11:12
Data OK
TIBCO(R) RTView(R) Display Server - TIBCO(R) RTView(R) for TIBCO BusinessEvents(R)
Version: 3.6.0.0 ALPHA
Configuration: TBE.3.6.0.0_20161201_000.23271-alpha_111
Build Number: 000.23271
Detailed Version Info For All Connected RTView Apps
Available Data Sources:
Alert
Cache Datasource
JMX
RtvAgent
SQL
XML
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

APPENDIX A Monitor Scripts

This section describes Monitor scripts and the **rtvservers.dat** configuration file. This section includes:

- Scripts
- rtvservers.dat

Scripts

This section describes scripts that are available for the Monitor.

The following scripts are available when used from an initialized command window. The scripts can be executed from a Windows Command Prompt or UNIX terminal window. On Windows, you can type the commands as described in this section. On UNIX systems, you must add **.sh** to each command. For example, **rtvapm_init.sh**.

These instructions assume use of a BASH or a BASH-compliant shell. The following scripts are available when used from an initialized command window. The scripts can be executed from a Windows Command Prompt or UNIX terminal window. On Windows, you can type the commands as described in this section. On UNIX systems, you must add **.sh** to each command. For example, **rtvapm_init.sh**. Also on UNIX systems, it is a requirement that the installation directory path not contain spaces.

These instructions assume use of a BASH or a BASH-compliant shell.

Script Name	Description
my_alert_actions.bat/sh	Sample script to define actions for alerts. Location: The project directory. Format: **my_alert_actions** (Append **.sh** on UNIX)
rtv_setup.bat/sh	Initializes a command prompt or terminal window. Location: **<installation directory>/bin** This script must be executed in the directory in which it resides. Format: **rtv_setup** (Append **.sh** on UNIX)
rtvapm_init.bat/sh	Initializes a command window. Location: **rtvapm** This script must be executed in the directory in which it resides. Format: **rtvapm_init** (Append **.sh** on UNIX)

start_cmd.bat	Starts an initialized Command Prompt window on Windows. Location: **<installation directory>/bin** This script must be executed in the directory in which it resides. You can also execute the script by double-clicking in an Explorer window.

start_rtv.bat/sh	Starts processes in an RTView configuration as specified in the **rtvservers.dat** configuration file.
	Location:
	rtvapm/common/bin
	This script must be executed in the project directory (the directory containing the **rtvservers.dat** file). This script requires **rtvapm_init.bat/sh** be executed first.
	An RTView configuration might include a Data Server or Display Server, an Historian and a Central Server Database. **start_rtv** only attempts to start processes it detects are not running. The action can be applied to all RTView configurations, a single RTView configuration or a single process in an RTView configuration.
	Before starting an RTView server, this script detects port conflicts caused by another server. If the conflict is caused by another RTView server, it returns a message identifying that server by its **rtvapm**. For example:
	`...start_rtv.bat: another dataserver running with JMX port 3268 under C:\rtview\RTViewDataServer\rtvapm`
	If the port conflict is caused by a non-RTView process, it returns a message similar to this, for example:
	`...start_rtv.bat: JMX port 3268 in use by PID 1234`
	In both cases the script includes this advice:
	`Warning: server not started, port conflict`

	To avoid port conflicts, run your start script with the **-portprefix:** command line argument to change the first two (**2**) digits of all your server ports. To persist these port changes, change the port prefix in the RTView Configuration Application or use the **-saveportprefix** command line argument. Additional arguments can be included on the command line in which case they are passed to every server specified by the command. Additional arguments can also be included in the **rtvservers.dat** file, in which case they are only applied to the specific server in whose command they are included. **Note:** If you use the **-properties** or **-propfilter** argument with **start_rtv**, you should also use them with **status_rtv** and **stop_rtv**. Those commands use the JMX ports defined for the server, and if any of the properties specified by **-properties** or **-propfilter** arguments change those ports, subsequent commands will be unable to find the server unless also given those properties.
	–console (or **–c**) - Start the processes with a command window (which is useful for testing).

	When used without arguments, this script returns usage information and a list of available configurations. For example, **start_rtv** returns:
	`Usage: start_rtv config or 'all' [server or 'all'] [args...]`
	`Available configs:`
	`default`
	`dataserver`
	`historian`
	`displayserver`
	`database`
	`sender`
	`dataserver`

	all Starts all RTView configurations that are specified in the **rtvservers.dat** file. **all** applies the action to all RTView configurations specified in the **rtvservers.dat** file (and corresponding servers or clients specified in each configuration). **Note:** When multiple configurations are specified in the **rtvservers.dat** file and they have different project settings directory locations, the **all** argument processes all the configurations. However, if the configurations have the same project settings directory locations, the **all** argument processes only the first configuration as the others are considered alternative configurations. Example: **start_rtv all** **(Append .sh on UNIX)**
	[Configuration Name] Starts a single RTView configuration specified in the **rtvservers.dat** file: **start_rtv [Configuration Name]** **(Append .sh on UNIX)** Configuration Name is the RTView configuration name specified in the **rtvservers.dat** file. The action applies to all servers or clients specified in the configuration. Example: **start_rtv web_deployment** **(Append .sh on UNIX)**
	[Server Name] Starts a single process in an RTView configuration specified in the **rtvservers.dat** file: **start_rtv [Configuration Name] [Server Name]** **(Append .sh on UNIX)** Server Name is the name of a server or client member in the configuration. For example, **dataserver**, **displayserver**, **historian** and **database**. The action applies only to that server or client in the configuration. Example: **start_rtv web_deployment dataserver**

	(Append .sh on UNIX)
	Use With Secured JMX Ports This script works with RTView servers whose JMX ports are secured with either a username and password, or with SSL. You provide the scripts with the necessary credential information and the scripts manage authentication with the server. There are two ways that you can provide credential information to the scripts: via command-line arguments and via properties placed in any property file that is used by the server. **Securing with username and password** • To secure with a username and password via command-line, use the arguments as follows: **-jmxuser:...** **-jmxpass:...** • To secure with a username and password in a property file, use the properties as follows: **sl.rtview.jmxremote.username=...** **sl.rtview.jmxremote.password=....** **Securing with SSL** To secure with SSL, you provide the client KeyStore and TrustStore locations and their corresponding passwords. • To secure with SSL via command-line, use the arguments as follows: **-sslkeystore:...** **-sslkeystorepass:...** **-ssltruststore:...** **-ssltruststorepass:...** • To secure with SSL in a property file, use the properties as follows: **sl.rtview.ssl.client.keyStore=...** **sl.rtview.ssl.client.keyStorePassword=...** **sl.rtview.ssl.client.trustStore=...** **sl.rtview.ssl.client.trustStorePassword=....** **Password Encryption** To encrypt the passwords in your properties files, use the command-line tool "encode_string", for example: **encode_string encoder2 password** This will give you an encrypted value for "password" that you can use in your properties.
start_server.bat/sh	Starts the RTView DataServer. Location: **<installation directory>** This script must be executed in the directory in which it resides. You can also execute the script by double-clicking in an Explorer window. Format: **start_server** (Append **.sh** on UNIX)
start_servers.bat/sh	Starts the RTViewCentral servers.

	Location: **<installation directory>/bin** This script must be executed in the directory in which it resides. You can also execute the script by double-clicking in an Explorer window. Format: **start_servers** (Append **.sh** on UNIX)
start_tomcat.bat/sh	Starts Apache Tomcat. Location: **<installation directory>/bin** This script must be executed in the directory in which it resides. You can also execute the script by double-clicking in an Explorer window. Format: **start_tomcat** (Append **.sh** on UNIX)
status_rtv.bat/sh	Returns the status of all RTView configurations that are specified in the **rtvservers.dat** configuration file. Location: **rtvapm/common/bin** This script must be executed in the project directory (the directory containing the **rtvservers.dat** file). This script requires **rtvapm_init.bat/sh** be executed first. This action uses defined JMX ports. An RTView configuration might include a Data Server, a Display Server or Viewer, an Historian and a Central Server Database. **status_rtv** only attempts to start processes it detects are not running. The action can be applied to all RTView configurations, a single RTView configuration or a single process in an RTView configuration. Additional arguments can be included on the command line in which case they are passed to every server specified by the command. Additional arguments can also be included in the **rtvservers.dat** file, in which case they are only applied to the specific server in whose command they are included. Note that if you use **-properties** or **-propfilter** arguments with **start_rtv**, you should also use them with **status_rtv** and **stop_rtv**. Those commands use the JMX ports defined for the server, and if any of the properties specified by **-properties** or **-propfilter** arguments change those ports, subsequent commands will be unable to find the server unless also given those properties.
	all Returns the status of all RTView configurations specified in the **rtvservers.dat** file. **Note:** When multiple configurations are specified in the **rtvservers.dat** file and they have different project settings directory locations, the **all** argument processes all the configurations. However, if the configurations have the same project settings directory locations, the **all** argument processes only the first configuration as the others are considered alternative configurations. Example: **status_rtv all** (Append **.sh** on UNIX)
	[Configuration Name] Returns the status of a single RTView configuration specified in the **rtvservers.dat** file:

	status_rtv [Configuration Name] **(Append .sh on UNIX)** Configuration Name is the RTView configuration name specified in the **rtvservers.dat** file. The action applies to all servers or clients specified in the configuration. Example: **status_rtv web_deployment** (Append **.sh** on UNIX)
	[Server Name] Returns the status of a single process in an RTView configuration specified in the **rtvservers.dat** file: **status_rtv [Configuration Name] [Server Name]** **(Append .sh on UNIX)** Server Name is the name of a server or client member in the configuration. For example, **dataserver**, **displayserver**, **historian** and **database**. The action applies only to that server or client in the configuration. Example: **status_rtv web_deployment dataserver** **(Append .sh on UNIX)**
	Use With Secured JMX Ports This script works with RTView servers whose JMX ports are secured with either a username and password, or with SSL. You provide the scripts with the necessary credential information and the scripts manage authentication with the server. There are two ways that you can provide credential information to the scripts: via command-line arguments and via properties placed in any property file that is used by the server. **Securing with username and password** • To secure with a username and password via command-line, use the arguments as follows: **-jmxuser:...** **-jmxpass:...** • To secure with a username and password in a property file, use the properties as follows: **sl.rtview.jmxremote.username=...** **sl.rtview.jmxremote.password=....** **Securing with SSL** To secure with SSL, you provide the client KeyStore and TrustStore locations and their corresponding passwords. • To secure with SSL via command-line, use the arguments as follows: **-sslkeystore:...** **-sslkeystorepass:...** **-ssltruststore:...** **-ssltruststorepass:...** • To secure with SSL in a property file, use the properties as follows: **sl.rtview.ssl.client.keyStore=...** **sl.rtview.ssl.client.keyStorePassword=...** **sl.rtview.ssl.client.trustStore=...** **sl.rtview.ssl.client.trustStorePassword=....**

	Password Encryption To encrypt the passwords in your properties files, use the command-line tool "encode_string", for example: **encode_string encoder2 password** This will give you an encrypted value for "password" that you can use in your properties.
status_server.bat/sh	Returns the status of the RTView DataServer. Location: **<installation directory>** This script must be executed in the project directory (the directory containing the **rtvservers.dat** file). Format: **status_server** (Append **.sh** on UNIX)
status_servers.bat/sh	Returns the status of the RTViewCentral servers (as well as the Solace PubSub+ Monitor in RTViewSolaceMonitor). Location: **<installation directory>/bin** This script must be executed in the project directory (the directory containing the **rtvservers.dat** file). Format: **status_servers** (Append **.sh** on UNIX)
stop_rtv.bat/sh	Stops processes in an RTView configuration as specified in the **rtvservers.dat** configuration file. Location: **rtvapm/common/bin** This script must be executed in the project directory (the directory containing the **rtvservers.dat** file). This script requires **rtvapm_init.bat**/**sh** be executed first. This action uses defined JMX ports. An RTView configuration might include a Data Server or a Display Server, an Historian and a Central Server Database. **stop_rtv** only attempts to start processes it detects are not running. The action can be applied to all RTView configurations, a single RTView configuration or a single process in an RTView configuration. Additional arguments can be included on the command line in which case they are passed to every server specified by the command. Additional arguments can also be included in the **rtvservers.dat** file, in which case they are only applied to the specific server in whose command they are included. Note that if you use **-properties** or **-propfilter** arguments with **start_rtv**, you should also use them with **status_rtv** and **stop_rtv**. Those commands use the JMX ports defined for the server, and if any of the properties specified by **-properties** or **-propfilter** arguments change those ports, subsequent commands will be unable to find the server unless also given those properties. Location: **project directory** **This script must be executed in the project directory (the directory containing the rtvservers.dat file).**
	all Stops all RTView configurations that are specified in the **rtvservers.dat** file. **all** applies the action to all RTView

	configurations specified in the **rtvservers.dat** file (and corresponding servers or clients specified in each configuration). **Note:** When multiple configurations are specified in the **rtvservers.dat** file and they have different project settings directory locations, the **all** argument processes all the configurations. However, if the configurations have the same project settings directory locations, the **all** argument processes only the first configuration as the others are considered alternative configurations. Example: **stop_rtv all** **(Append .sh on UNIX)**
	[Configuration Name] Stops a single RTView configuration specified in the **rtvservers.dat** file: **stop_rtv [Configuration Name]** **(Append .sh on UNIX)** Configuration Name is the RTView configuration name specified in the **rtvservers.dat** file. The action applies to all servers or clients specified in the configuration. Example: **stop_rtv web_deployment** **(Append .sh on UNIX)**
	[Server Name] Stops a single process in an RTView configuration specified in the **rtvservers.dat** file: **stop_rtv [Configuration Name] [Server Name]** **(Append .sh on UNIX)** **Server Name** is the name of a server or client member in the configuration. For example, **dataserver**, **displayserver**, **historian** and **database**. The action applies only to that server or client in the configuration. Example: **stop_rtv web_deployment dataserver** **(Append .sh on UNIX)**
	Use With Secured JMX Ports This script works with RTView servers whose JMX ports are secured with either a username and password, or with SSL. You provide the scripts with the necessary credential information and the scripts manage authentication with the server. There are two ways that you can provide credential information to the scripts: via command-line arguments and via properties placed in any property file that is used by the server. **Securing with username and password** • To secure with a username and password via command-line, use the arguments as follows: **-jmxuser:...** **-jmxpass:...** • To secure with a username and password in a property file, use the properties as follows: **sl.rtview.jmxremote.username=...** **sl.rtview.jmxremote.password=....** **Securing with SSL** To secure with SSL, you provide the client KeyStore and

	TrustStore locations and their corresponding passwords. • To secure with SSL via command-line, use the arguments as follows: **-sslkeystore:...** **-sslkeystorepass:...** **-ssltruststore:...** **-ssltruststorepass:...** • To secure with SSL in a property file, use the properties as follows: **sl.rtview.ssl.client.keyStore=...** **sl.rtview.ssl.client.keyStorePassword=...** **sl.rtview.ssl.client.trustStore=...** **sl.rtview.ssl.client.trustStorePassword=....** **Password Encryption** To encrypt the passwords in your properties files, use the command-line tool "encode_string", for example: **encode_string encoder2 password** This will give you an encrypted value for "password" that you can use in your properties.
stop_server.bat/sh	Stops the RTView DataServer. Location: **<installation directory>** This script must be executed in the directory in which it resides. Format: **stop_server** (Append **.sh** on UNIX)
stop_servers.bat/sh	Stops the RTViewCentral servers. Location: **<installation directory>/bin** This script must be executed in the directory in which it resides. You can also execute the script by double-clicking in an Explorer window. Format: **stop_servers** (Append **.sh** on UNIX)
stop_tomcat.bat/sh	Stops Apache Tomcat. Location: **<installation directory>/bin** This script must be executed in the directory in which it resides. Format: **start_tomcat** (Append **.sh** on UNIX)
update_wars.bat/sh	Creates/updates the primary Monitor servlets. Location: **<installation directory>/projects/rtview-server** This script must be executed in the directory in which it resides. This script requires **rtvapm_init.bat**/**sh** be executed first.

	Format: **update_wars.sh [appname [host [portprefix]]]** For example: **update_wars.sh my-appname my-hostname 99** The name, host, and portprefix are declared in variables at the top of the script for easy editing, and can be passed into the scripts on the command-line. **-secure** Use the "-secure" argument to update the rtvquery war with security enabled. You can use **?** or **help** to get a usage message. For example: **update_wars.sh help** You can edit other variables at the top of the scripts to set properties for high-availability (HA). **Set HA_HOST** to the hostname of the backup data server. **Set HA_DISPLAYHOST** to the hostname of the backup display server. **Set HA_FAILBACK** to true to automatically reconnect to the primary display server.

rtvservers.dat

This section describes the **rtvservers.dat** configuration file which is used to manage your TIBCO® RTView® for TIBCO BusinessEvents® deployment and RTView processes. This section includes:

- Single Configuration File
- Multiple Configuration File

The **rtvservers.dat** text file contains one or more RTView configurations. An RTView configuration is a group of servers that should be started together. For example, the configuration might include any of the following: a Data Server, Historian, HSQLDB database, and a Display Server (for a Web Deployment). The **rtvservers.dat** file is used when the following scripts are executed:

- start_rtv Starts RTView processes specified in the **rtvservers.dat** file.
- stop_rtv Stops the RTView processes specified in the **rtvservers.dat** file.
- status_rtv Returns status information for RTView processes specified in the **rtvservers.dat** file.

Single Configuration File

The following **rtvservers.dat** file, located in your project directory, contains a single RTView configuration, named **default**.

```
default . dataserver rundata
default . historian runhist -ds
```

default . displayserver rundisp -ds

default . database rundb

Note: The last line in the **rtvservers.dat** file must end with a new line, or be followed by a blank line.

In this example, to start the **default** configuration type: **start_rtv default** or **start_rtv all**. To start a single server in the configuration, type **start_rtv <Configuration Name> <Server Name>**. For example: **start_rtv default displayserver**.

Each line has the following format consisting of four fields:

<Configuration Name> <Project Settings Directory Location> <Property Filter Identifying the Server> <Command>

<Configuration Name>	The name of the RTView configuration (**default** in this example).
<Project Settings Directory Location>	The RTView project settings directory location, relative to the location of the **rtvservers.dat** file (**.**, the current directory, in this example).
<Property Filter Identifying the Server>	The property filter that identifies the server, which is the property filter under which the server's JMX port is defined. By default, this is the server name, such as **dataserver**, **displayserver** and **historian**.
<Command>	The script used to start the process. Valid values are: • rundata: Starts the Data Server. • runhist: Starts the Historian. • rundisp: Starts the Display Server. • rundb: Starts the HSQLDB Database.

Multiple Configuration File

When multiple configurations are specified in the **rtvservers.dat** file and they have different project settings directory locations, the **all** argument processes all the configurations. However, if the configurations have the same project settings directory locations, the all argument processes only the first configuration as the others are considered alternative configurations. Alternative configurations allow you to alternate between two configurations for a single TIBCO® RTView® for TIBCO BusinessEvents® deployment.

For example, the following **rtvservers.dat** file contains two configurations, **bemon** and **emsmon**. Note that the project settings directory locations differ (**./bemon** and **./emsmon**, respectively).

bemon ./bemon dataserver rundata

bemon ./bemon historian runhist -ds

bemon ./bemon displayserver rundisp -ds

emsmon ./emsmon dataserver rundata

emsmon ./emsmon historian runhist -ds

emsmon ./emsmon displayserver rundisp –ds

Because the project settings directory locations differ, you can use type **start_rtv all** to start both configurations. To start only the bwmon configuration, type: **start_rtv bemon**. To start

a single server in the **bemon** configuration, type **start_rtv <Configuration Name> <Server Name>**. For example: **start_rtv bemon displayserver**.

APPENDIX B Alert Definitions

This section describes alerts for TIBCO BusinessEvents and their default settings.

TbeBackingStoreEraseRateHigh	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 **80** and the alarm default threshold is **95**.
TbeBackingStoreLoadRateHigh	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 **80** and the alarm default threshold is **95**.
TbeBackingStoreStoreRateHigh	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 **80** and the alarm default threshold is **95**.
TbeClusterMalformed	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. 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".
TbeDestinationStatusRecvdEventsRateHigh	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 **80** and the alarm default threshold is **95**.
TbeNodeConceptsGetRateHigh	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 **80** and the alarm default threshold is **95**.
TbeNodeConceptsPutRateHigh	This alert executes a single warning alert and a single alarm alert if the rate at which concepts are written to the cache exceeds the specified threshold. The warning default threshold is **80** and the alarm default threshold is **95**.
TbeNodeConceptsRemoveRateHigh	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 **80** and the alarm default threshold is **95**.
TbeNodeConnectionLoss	This discrete alert executes when the JMX Connection to the TIBCO BusinessEvents agent is lost (the TCP connection flag for an engine is **false**).
TbeNodeEventsGetRateHigh	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 **80** and the alarm default threshold is **95**.
TbeNodeEventsPutRateHigh	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 **80** and the alarm default threshold is **95**.
TbeNodeEventsRemoveRateHigh	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 **80** and the alarm default threshold is **95**.
TbeObjectTableExtIdSize	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 **9000** and the alarm default threshold is **10000**.
TbeObjectTableSize	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 **9000** and the alarm default threshold is **10000**.
TbeRuleFiringRateHigh	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 **80** and the alarm default threshold is **95**.

JvmCpuPercentHigh	This alert executes a single warning alert and a single alarm alert if the percent CPU used by the JVM exceeds the limit defined for that JVM. The warning default threshold is **50** and the alarm default threshold is **75**.
JvmGcDutyCycleHigh	This alert executes a single warning alert and a single alarm alert if the percent memory used for duty cycling by the JVM exceeds the limit defined for that JVM. The warning default threshold is **50** and the alarm default threshold is **75**.
JvmMemoryUsedHigh	This alert executes a single warning alert and a single alarm alert if the percent heap memory used by the JVM, relative to the maximum heap memory available, exceeds the limit defined for that JVM. The warning default threshold is **50** and the alarm default threshold is **75**.
JvmNotConnected	This discrete alert executes when the JVM Connection is lost (the connection flag for an engine is **false**). This is an indication that the JVM might have stopped functioning.
JvmStaleData	This alert executes when the metrics for a given Java virtual machine expire, which means that the Data Server is unable to collect new data from that JVM. This could be accompanied by a **JvmNotConnected** alert, if enabled. The expiration time is set by the **$jvmRowExpirationTime** property, which defaults to seconds.

APPENDIX C RTView Configuration Application

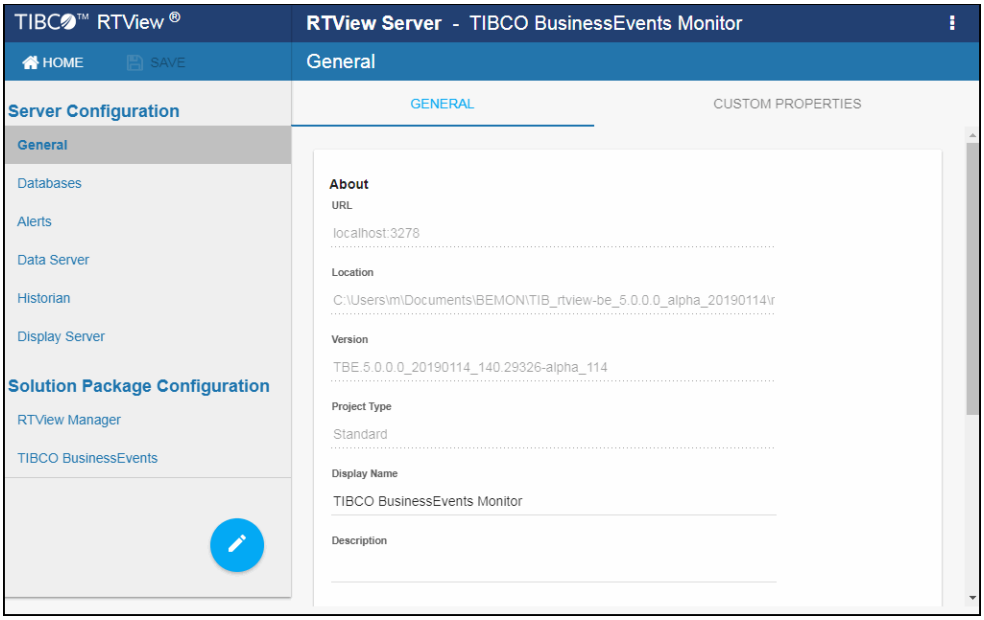

The RTView Configuration Application is a tool that you can use to help configure the Monitor by defining various properties and connections via an easy-to-use interface. The RTView Configuration Application consists of three different sections: the top-level **Projects** page, the **Server Configuration** view, and the **Solution Package Configuration** view. This section will provide high-level definitions for each option within each view. More detailed descriptions on how this tool can be used to set up the Monitor can be found in the Configuration chapter, as well as in the Quick Start chapter.

This section contains the following:

- Accessing the RTView Configuration Application
- Projects Page
- Server Configuration View
- Solution Package Configuration View

Accessing the RTView Configuration Application

There are two ways you can access the RTView Configuration Application:

- Via URL
- Via a Button in the Monitor (if "Login" is enabled)

Via URL

You can access the RTView Configuration Application via URL by performing the following steps:

1. Download and extract the TBEMON compressed .zip file..
2. Set the **JAVA_HOME** environment variable.
3. Run **start_server** from the **TIB_rtview-be** directory to start all servers.
4. Open a browser and enter **http://localhost:3270/rtview-tbemon-rtvadmin**.

See Quick Start for additional details.

> **Note:** Once you have finished making changes in the RTView Configuration Application, you must restart your data server for your changes to take place in the Monitor.

Via a Button in the Monitor

> **Note:** "Login" must be enabled and you must log in with administrator privileges in order for the RTView Configuration Application button to display in the Monitor. See Enabling Login in the Monitor for more information. The RTView Configuration Application button is only available in the Display Server version of the User Interface. The button will not display in the HTML User Interface.

You can access the RTView Configuration Application via a button in the Monitor by performing the following steps:

1. Download and extract the TBEMON compressed .zip file..
2. Set the **JAVA_HOME** environment variable.
3. Run **start_server** from your **TIB_rtview-be** directory to start all servers.
4. Open a browser and enter **http://localhost:3270/rtview-tbemon-classic**.
   The Monitor displays.
5. Click in the Monitor to open the RTView Configuration Application.

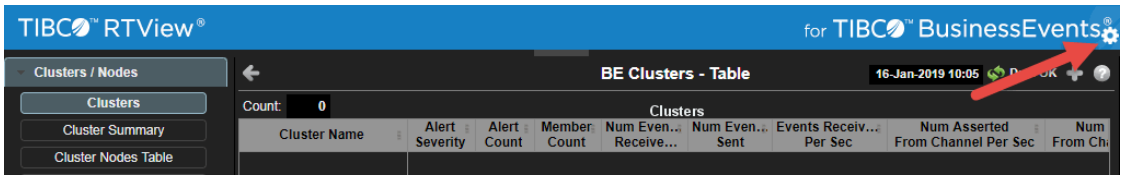

See Quick Start for additional details.

> **Note:** Once you have finished making changes in the RTView Configuration Application, you must restart your data server for your changes to take place in the Monitor.

Projects Page

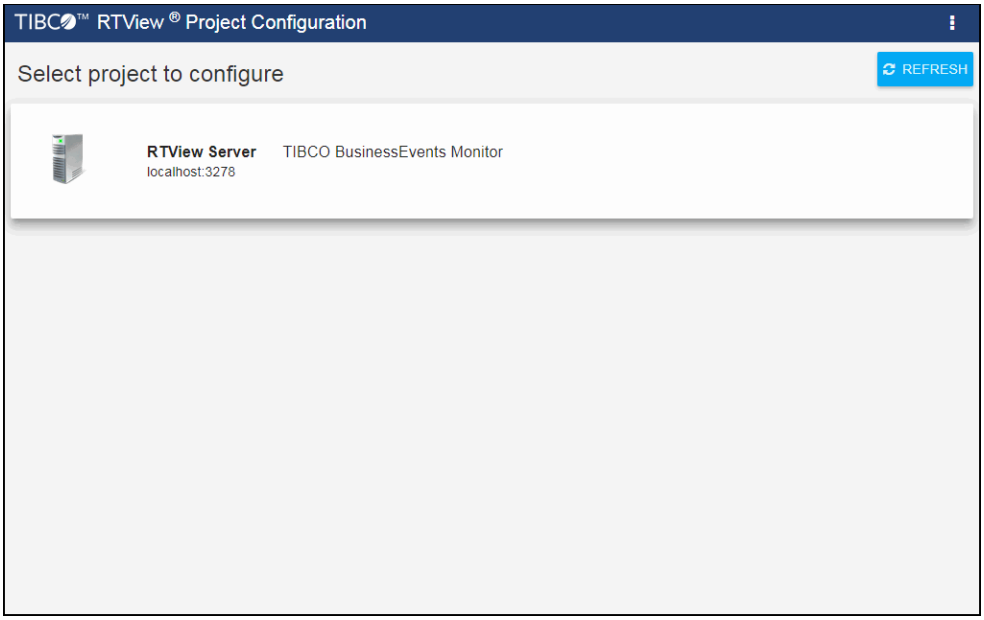

The Projects Page lists the project(s) in your project directory. Click the project to access the Configuration Views.

Server Configuration View

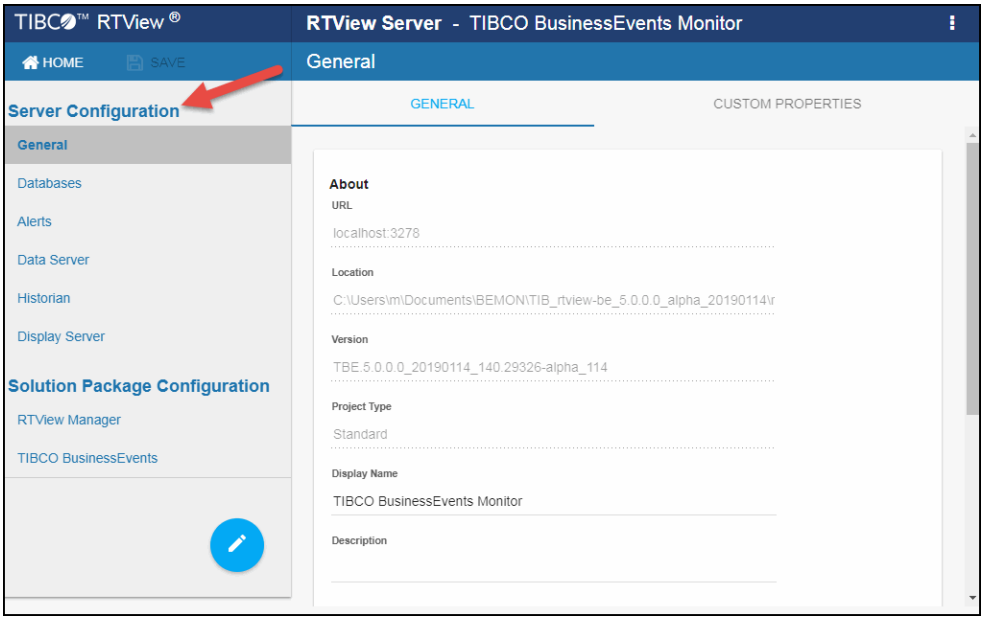

The **Server Configuration** View provides options that allow you to modify the default settings for the project including the project name and default port, define the alert threshold database connection and alert notification settings, define custom properties, define data

server properties, define display server properties, and define the historian database connection and other historian properties. This section contains the following:

- General
- Databases
- Alerts
- Data Server
- Historian
- Display Server

Tip: Gray text shows the default setting for the field which you can edit. To return to the default setting, delete the text you entered.

General

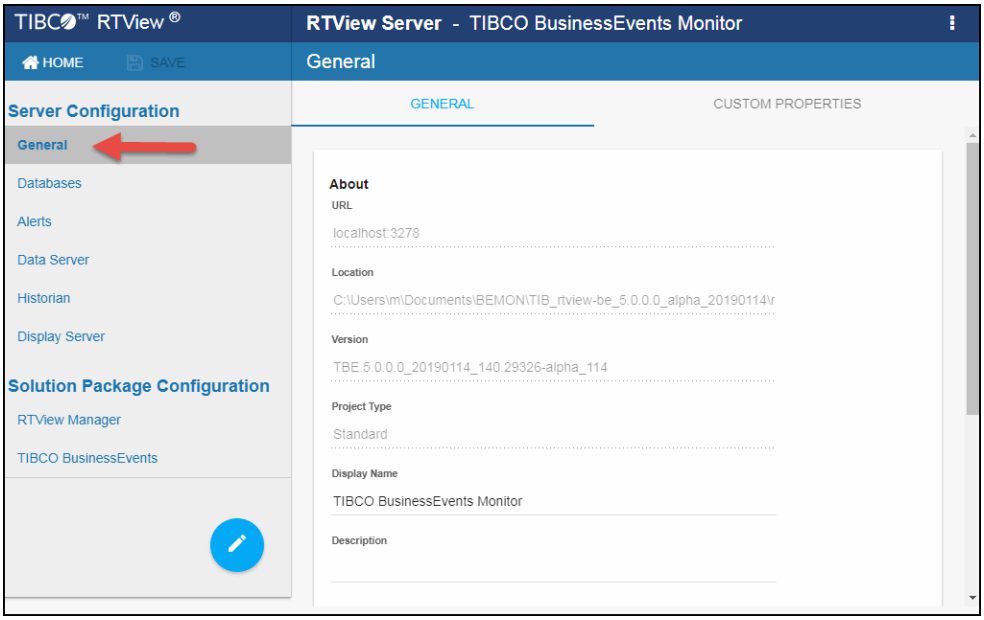

The **General** option consists of two different tabs that allow you to define the values for the project, specify the port, and define any custom properties you might need to create. The available tabs are:

- General Tab
- Custom Properties Tab

General Tab

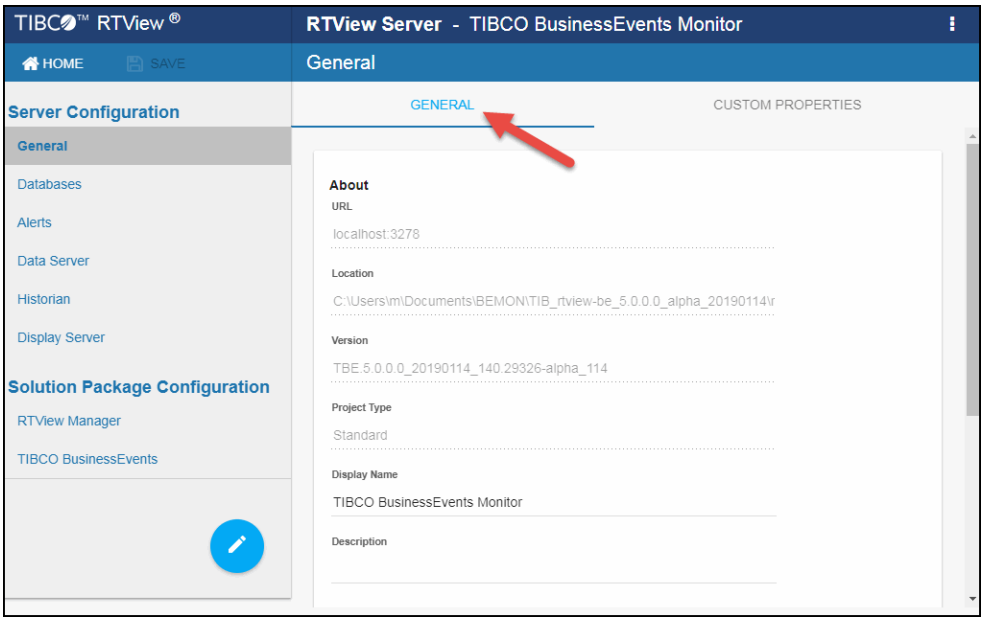

This tab contains the following regions:

About

URL: Displays the URL used to the connect to the server. This field cannot be edited.

Location: Displays the project directory location (path). This field cannot be edited.

Version: Displays the current version of TIBCO BusinessEvents installed. This field cannot be edited.

Project Type: Displays the type of project (Standard, Sender, or ConfigClient). This field cannot be edited.

Display Name: Displays the default name for the project and displays on the Home/**RTView Project Configuration** (top level) page. This field can be edited.

Description: Optionally specify a description that will display on the Home/**RTView Project Configuration** (top level) page.

Identifier

Project ID: Displays a default unique identifier for the project, which you can modify.

Ports

Port Prefix: Displays the default port prefix (first two numbers used for the port) that will be used for all ports, which you can modify. The latter two numbers in the port are predefined and cannot be modified. Click **Show Port Assignments** to view the Port Assignments.

Custom Properties Tab

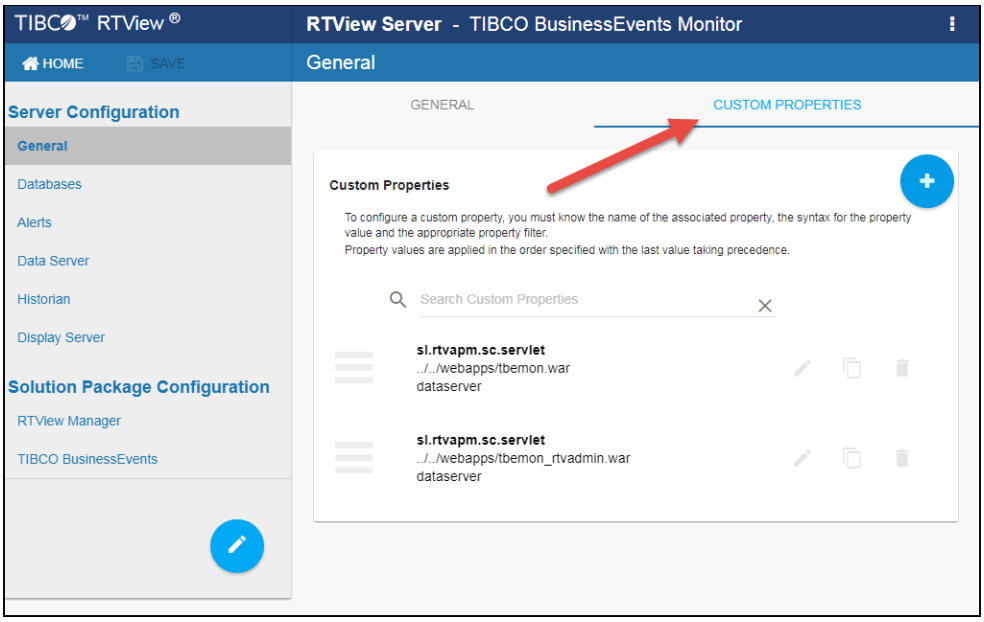

The **Custom Properties** tab allows you to create custom properties. Property values are applied in the order specified with the last value taking precedence. To create properties you need the name of the associated property, the syntax for the property value, and the appropriate property filter. Click the + icon to open the **Add Property** dialog, which has the following fields:

Name: (Required) The name of the associated property.

Value: (Optional) The value for the associated property (using the correct syntax).

Filter: (Optional) The filter for the associated property.

Comment: (Optional) Enter useful details about the property and its behavior for yourself and other users.

See Configuring Alert Notification for an example of when you can use **Custom Properties**.

Databases

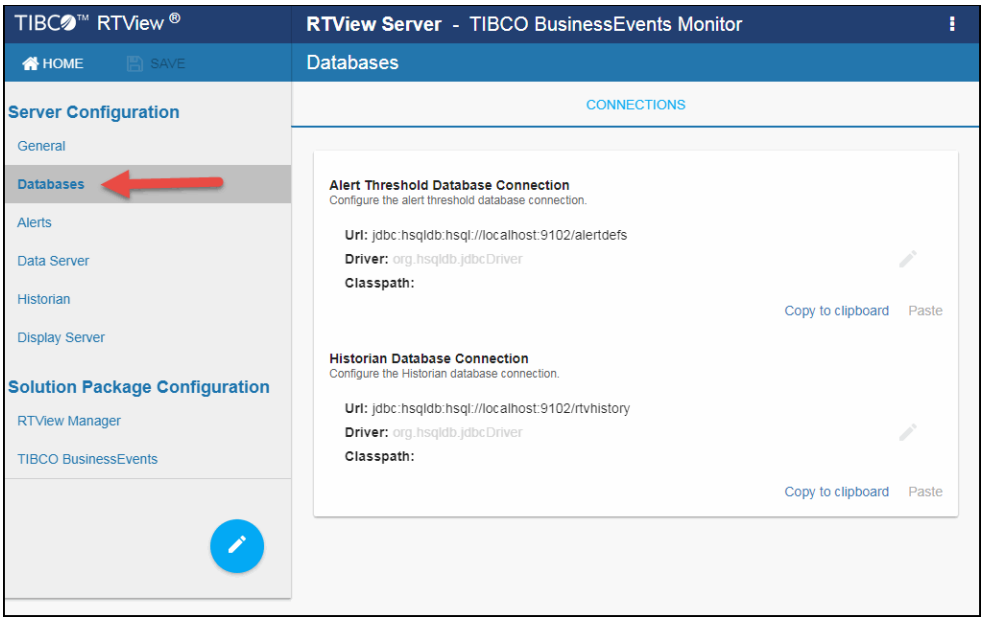

The **Databases** option consists of the **Connections** tab that allows you to define Alert Threshold Database and Historian Database connections.

- Connections Tab

Connections Tab

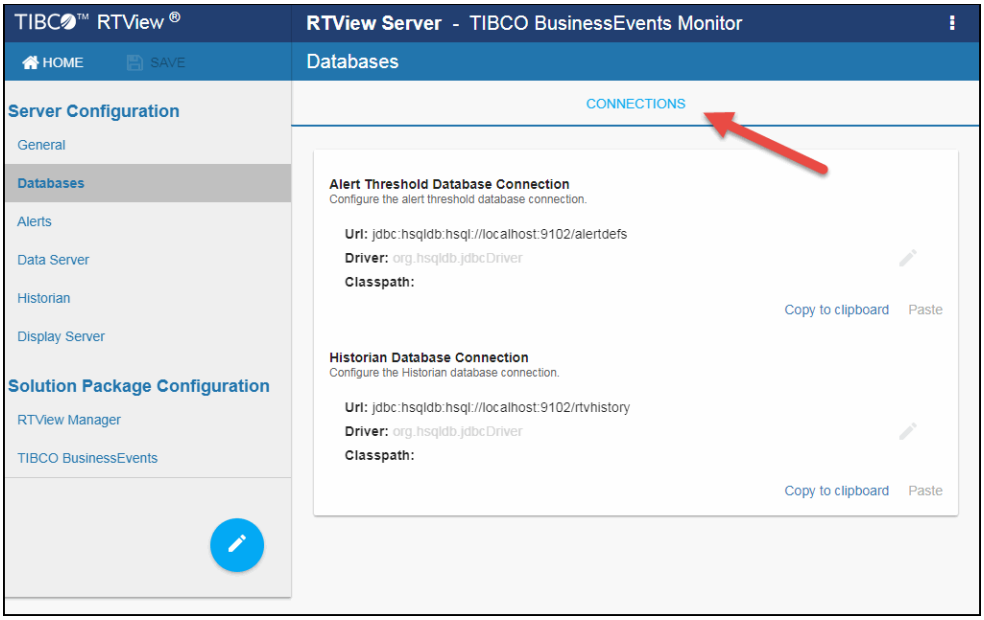

This tab contains the following regions:

Alert Threshold Database Connection

If you want to use local alert threshold database connection, add the connection information where:

URL: The complete URL for the database connection.

Driver: The full name for the driver.

Classpath: The complete classpath for the jar location.

Username: The username is used when creating the connection. This field is optional.

Password: This password is used when creating the connection. This field is optional. By default, the password entered is hidden. Click the icon to view the password text.

Run Queries Concurrently: When selected, database queries are run concurrently.

Historian Database Connection

URL: The complete URL for the database connection.

Driver: The full name for the driver.

Classpath: The complete classpath for the jar location.

Username: The username is used when creating the connection. This field is optional.

Password: This password is used when creating the connection. This field is optional. By default, the password entered is hidden. Click the icon to view the password text.

Run Queries Concurrently: When selected, database queries are run concurrently.

Alerts

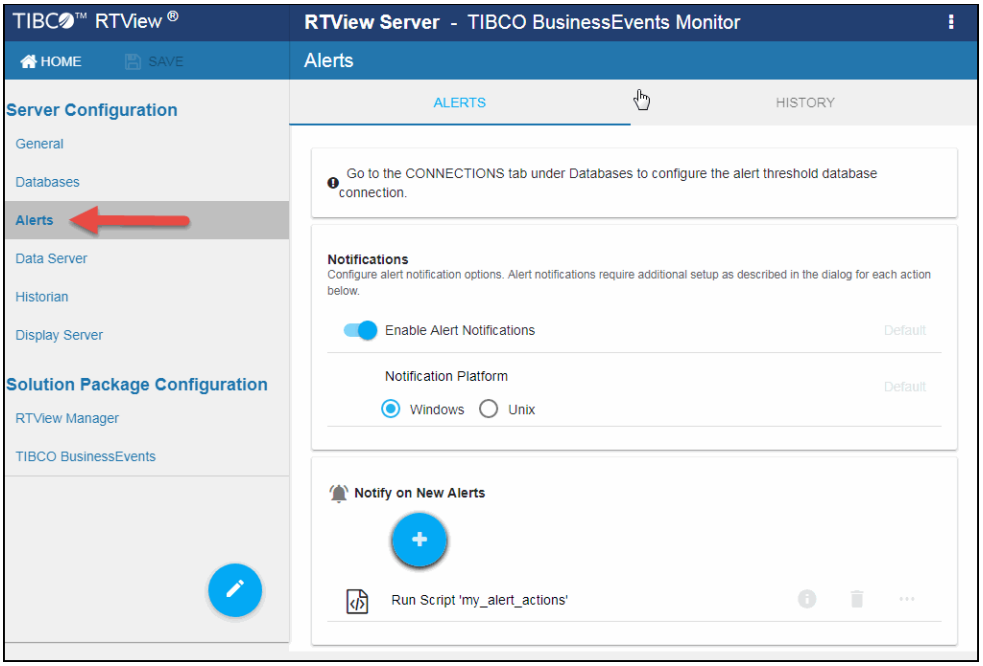

The Alerts option consists of the Alerts tab and the History tab, which allow you to define the alert and history properties. Alert and Historian database connections are set up using the Databases option. The following tabs are available:

- Alerts Tab
- History Tab

Alerts Tab

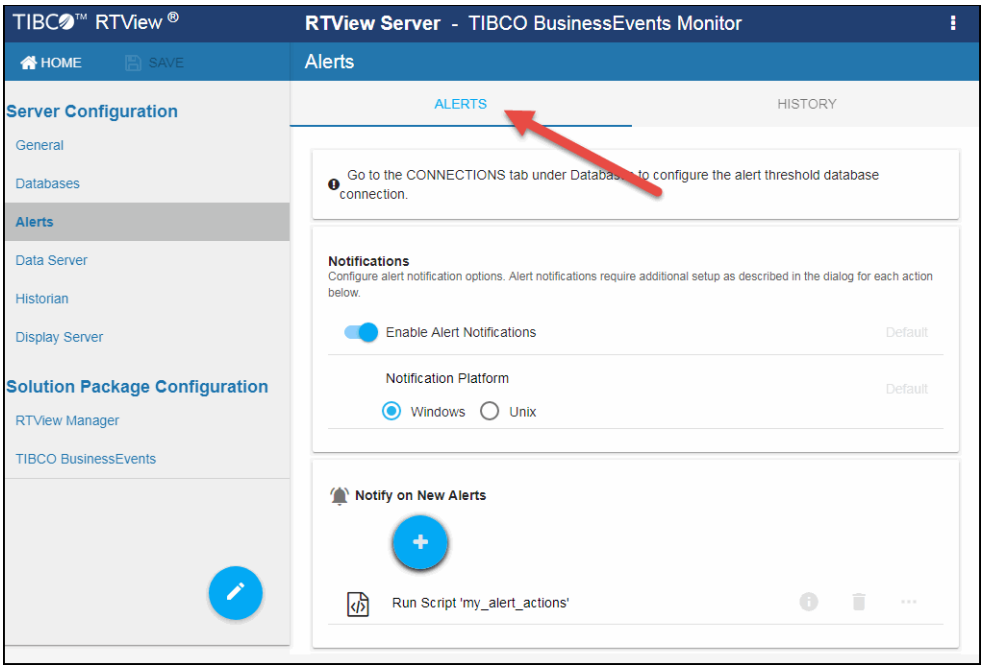

This tab contains the following regions:

Notifications

- **Enable Alert Notifications**: Selecting this toggle enables alert notifications to be sent.
- **Notification Platform**: Select the platform type (**UNIX** or **Windows**).

Alert Event Options

- **Notify on New Alerts**: A notification is executed every time a new alert is created.
- **Nofity on First Severity Change**: A notification is executed the first time the **Severity** changes for each alert.
- **Nofity on Cleared Alerts**: A notification is executed every time an alert is cleared.
- **Periodically Renotify on Unacknowledged Alerts**: Enter the **Renotification Interval** (number of seconds). A notification is executed for each unacknowledged alert per the interval you specify here. If the Renotification Interval is greater than **0** and no actions are defined, the **New Alerts** action will be used for renotifications.

Selecting the button next to each of the Alert Event Options displays the following options:

Run a Script

This alert notification action executes the following script in the **TIB_rtview-ems/projects/rtview-server** directory:

- **my_alert_actions.bat/sh** – New and First Severity Change
- **my_alert_actions.cleared.bat/sh** – Cleared
- **my_alert_actions.renotify.bat/sh** – Periodically Renotify

This action can only be added once per notification type. In addition to selecting this action in the Configuration Application, you must also modify the appropriate script to execute the actions for your notification. This script has access to the following fields from the alert: **Alert Name**, **Alert Index**, **ID**, **Alert Text** and **Severity**.

This alert notification action allows you to implement your alert notification actions using Java code. It executes the **my_alert_notification.$domainName.$alertNotifyType.$alertNotifyCol** command in your Custom Command Handler and passes the row from the alert table that corresponds to the alert.

Execute Java Code

This action can only be added once per notification type. In addition to selecting this action the Configuration Application you must also modify the custom command handler to execute the actions for your notification. A sample custom command handler is included under **projects/custom**. It prints the alert notification to the console. You will modify this command handler to implement your own notification actions.

Make the following entries:

- **Custom Command Handler Class Name**: Enter the fully qualified name of the Custom Command Handler class. This defaults to the sample Custom Command Handler in the **TIB_rtview-ems/projects/custom** directory.
- **Custom Command Handler Jar**: Enter the path and name of the jar containing the Custom Command Handler class. The path may be absolute or relative to the location of data server. This defaults to the sample Custom Command Handler in the **TIB_rtview-ems/projects/custom** directory.

Note that if you can only have one custom command handler per Data Server, so changing these settings for one notification event will change them for the rest of the notification events.

Send Email

This alert notification action sends an email. This action can be added multiple times per notification type. No additional setup is required beyond filling in the **Send Email** dialog in the Configuration Application.

Make the following entries:

- **SMTP Host**: The SMTP host address. This is required. Consult your administrator.
- **SMTP Port**: The SMTP port number. This is required. Consult your administrator.
- **From**: The email address to which to send the email. This is required.
- **To**: The email address to which to send the email. This is required and may contain multiple entries.
- **Subject**: The subject for the email. This is required. You can include the value from any column in the alert table in your subject. Click the **Show More** link at the bottom of the dialog

to see the alert column values you can use in the **Subject**.

- **Body**: The body of the email. This is optional. Click the **Show More** link at the bottom of the dialog to see the alert column values you can use in the **Subject**.
- **User**: The user name for the account from which you are sending the email. This is optional.
- **Password**: The password for the account from which you are sending the email. This is optional.

Send SNMP Trap

This alert notification action sends an SNMP Trap as described in **rtvapm/common/lib/SL-RTVIEW-EM-MIB.txt**. This action can be added multiple times per notification type. No additional setup is required beyond filling in the **Send Email** dialog in the Configuration Application

Make the following entries:

- **Trap Type**: Select the SNMP version of the trap. This is required.
- **Destination Address**: The system name or IP address of the receiving system. This is required.
- **Destination Port**: The UDP port on the receiving system. This is required.
- **Community Name**: (This field is visible when **Trap Type v2/v3** is selected.) The SNMP v2 Community Name string. This is required.

Run Command String

This alert notification action executes a specified command. This action can be added multiple times per notification type. Make the following entry:

Command String: Enter the command string for any command supported by RTView Classic. To enter a command string, you must know the correct syntax for the command. Contact Technical Support for assistance on syntax. You can include the value from any column in the alert table using the syntax in the Show More link at the bottom of the dialog.

Conditional Filter

This alert notification action alert allows you to execute different actions for different alerts based on information in the alert. For example, you can configure EMS alerts to send emails to your EMS team and Solace alerts to send emails to your Solace team. This action can be added multiple times per notification type.

To create a condition, make the following entries:

- **Alert Field**: Select an alert field: **Alert Name**, **Alert Index**, **Category, CI Name, Owner, Package, Primary Service** or **Severity**. This is required. Note that **CI Name** and **Primary Service** fields are for RTViewCentral only.
- **Operator**: Select one - **EQUALS**, **DOES NOT EQUAL**, **STARTS WITH**, **ENDS WITH** or **CONTAINS**. This is required.
- **Value**: Enter the value to which to compare the Alert Field.

Cannot contain wildcard characters. This is required.

- **Action(s)**: Select one or more actions to execute when this condition is met.

Persistence

Persist Alerts: When enabled, saves alerts to the database for high availability purposes.

History Tab

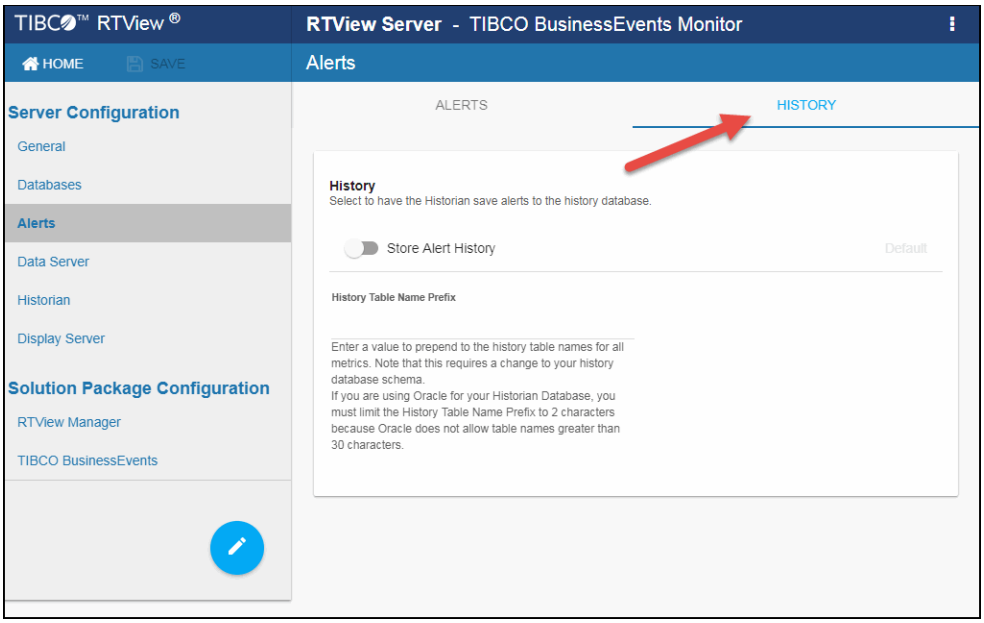

This tab contains the following region:

History

Store Alert History: Toggle to enable/disable **Store Alert History** to store alerts in the history database. This value is used in the **Alerts Table** (which makes it easier to find the alerts).

History Table Name Prefix: This field allows you to define a prefix that will be added to the database table names so that the Monitor can differentiate history data between data servers when you have multiple data servers with corresponding Historians using the same solution package(s) and database. In this case, each Historian needs to save to a different table, otherwise the corresponding data server will load metrics from both Historians on startup. Once you have defined the **History Table Name Prefix**, you will need to create the corresponding tables in your database as follows:

- Locate the .sql template for your database under **TIB_rtview-be/rtvapm/bemon/dbconfig** and make a copy of it
- Add the value you entered for the **History Table Name Prefix** to the beginning of all table names in the copied .sql template
- Use the copied .sql template to create the tables in your database

Security

All RTView processes (Data Server, Historian, Display Server) open JMX ports for monitoring which, by default, are not secured. The **Security** tab allows you secure these ports as well as specify credentials needed to connect to SSL secured servers from RTView's Solution Packages.

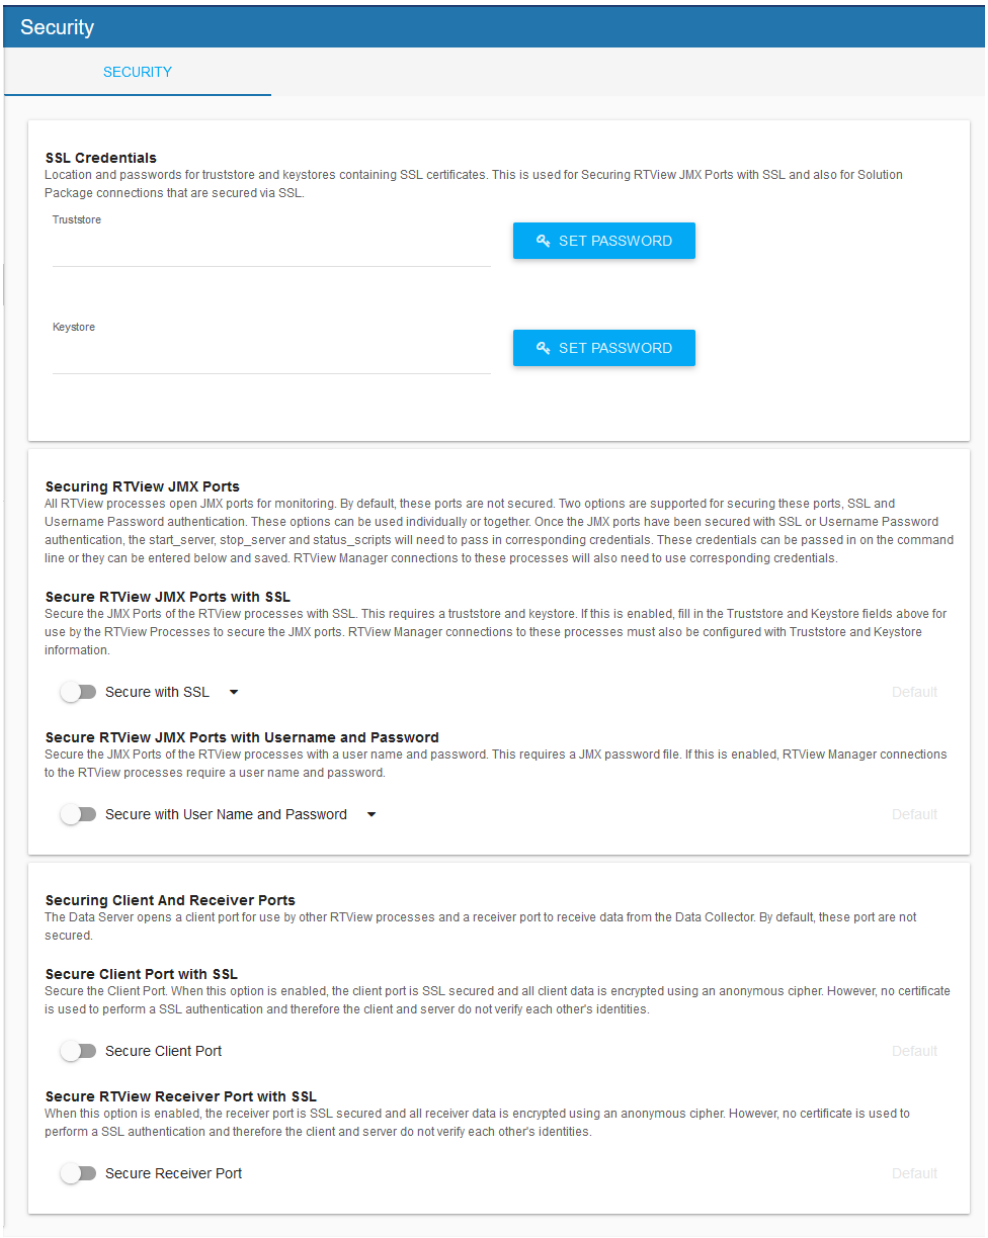

SSL Credentials

This region allows you to specify the path to the **Truststore** and **Keystore** files (and their associated passwords) that contain the SSL credentials needed to secure the RTView JMX Ports and/or access any SSL secured servers associated with RTView's Solution Packages. This is required if the **Secure with SSL** option is enabled (see below for details).

Optional: To obscure the credentials of the truststore and keystore in the output of the **ps** and **jps** commands, add the following custom property to each Data Server on which SSL Credentials have been configured:

```
Name: sl.rtview.jvm
Value: -Drtv.hidesslprops=true
Comment: hide ssl properties in ps/jps output
```

Securing RTView JMX Ports

This region provides a couple of options for securing the JMX ports that are opened by the RTView processes: **Secure with SSL** and/or **Secure with Username and Password**.

Secure with SSL

When toggled on, this option secures the JMX ports for the RTView processes with SSL. When the JMX ports are SSL secured, an SSL handshake is performed between the client and the server when the client attempts to connect. For this handshake, the client must provide a certificate the server trusts, and the server must provide a certificate the client trusts. A Keystore file contains the application's certificate and private key and a Truststore file contains the application's trusted certificates. These files are created by running the Java keytool on the command line. When this option is enabled, you must specify the path to the server's Truststore and Keystore files (and their associated passwords) in the **SSL Credentials** region (see above).

The **start_server**, **stop_server**, and **status_server** scripts are all connected to the JMX Ports of the RTView processes to execute operations and get status. If the JMX ports have been secured with SSL, these scripts need the path and passwords for the truststore and keystore files containing the client credentials in order to connect. You can either pass these in on the command line each time you run (**-sslkeystore:clientkeystore.jks-sslkeystorepass:clientkeystorepass-ssltruststore:clienttruststore.jks-ssltruststorepass:clienttruststorepass**) or you can fill in the fields under **SSL Credentials for RTView Scripts**.

The RTView Manager application also connects to the JMX Ports of the RTView processes in order to monitor them. If you are using the RTView Manager and the JMX ports have been secured with SSL, you must fill in the **SSL Credentials** on the **Security** tab of the RTView Manager Configuration Application to specify the path the truststore and keystore files containing the client credentials.

Secure with Username and Password

This region allows you to secure the JMX ports for RTView processes with a username/password. This can be used in addition to Securing with SSL (see above). If this option is enabled, you must specify the path to a JMX password file containing all valid user names and passwords.

Important! A JMX password file must be read-only to the owner. See Java documentation for details on the creation of a JMX remote password file.

The **start_server**, **stop_server**, and **status_server** scripts are all connected to the JMX Ports of the RTView processes to execute operations and get status. If the JMX ports have been secured with a username and password, the scripts need a valid user name and password in order to connect. You can either pass these into the command line each time you run (**-jmxuser:userName-jmxpass:myPassword**) or you can fill in the **Username and Password Credentials** and enable the **Use for Scripts** toggle.

The RTView Manager application also connects to the JMX Ports of the RTView processes in order to monitor them. If you are using the RTView Manager in RTViewCentral and the JMX ports have been secured with username and password, you must fill in the **Username and Password Credentials** that the RTView Manager should use to connect. If you are using the RTView Manager in a deliverable other than RTViewCentral, you will need to fill in the user name and password in the connection to this RTViewDataServer in the RTView Manager Configuration Application.

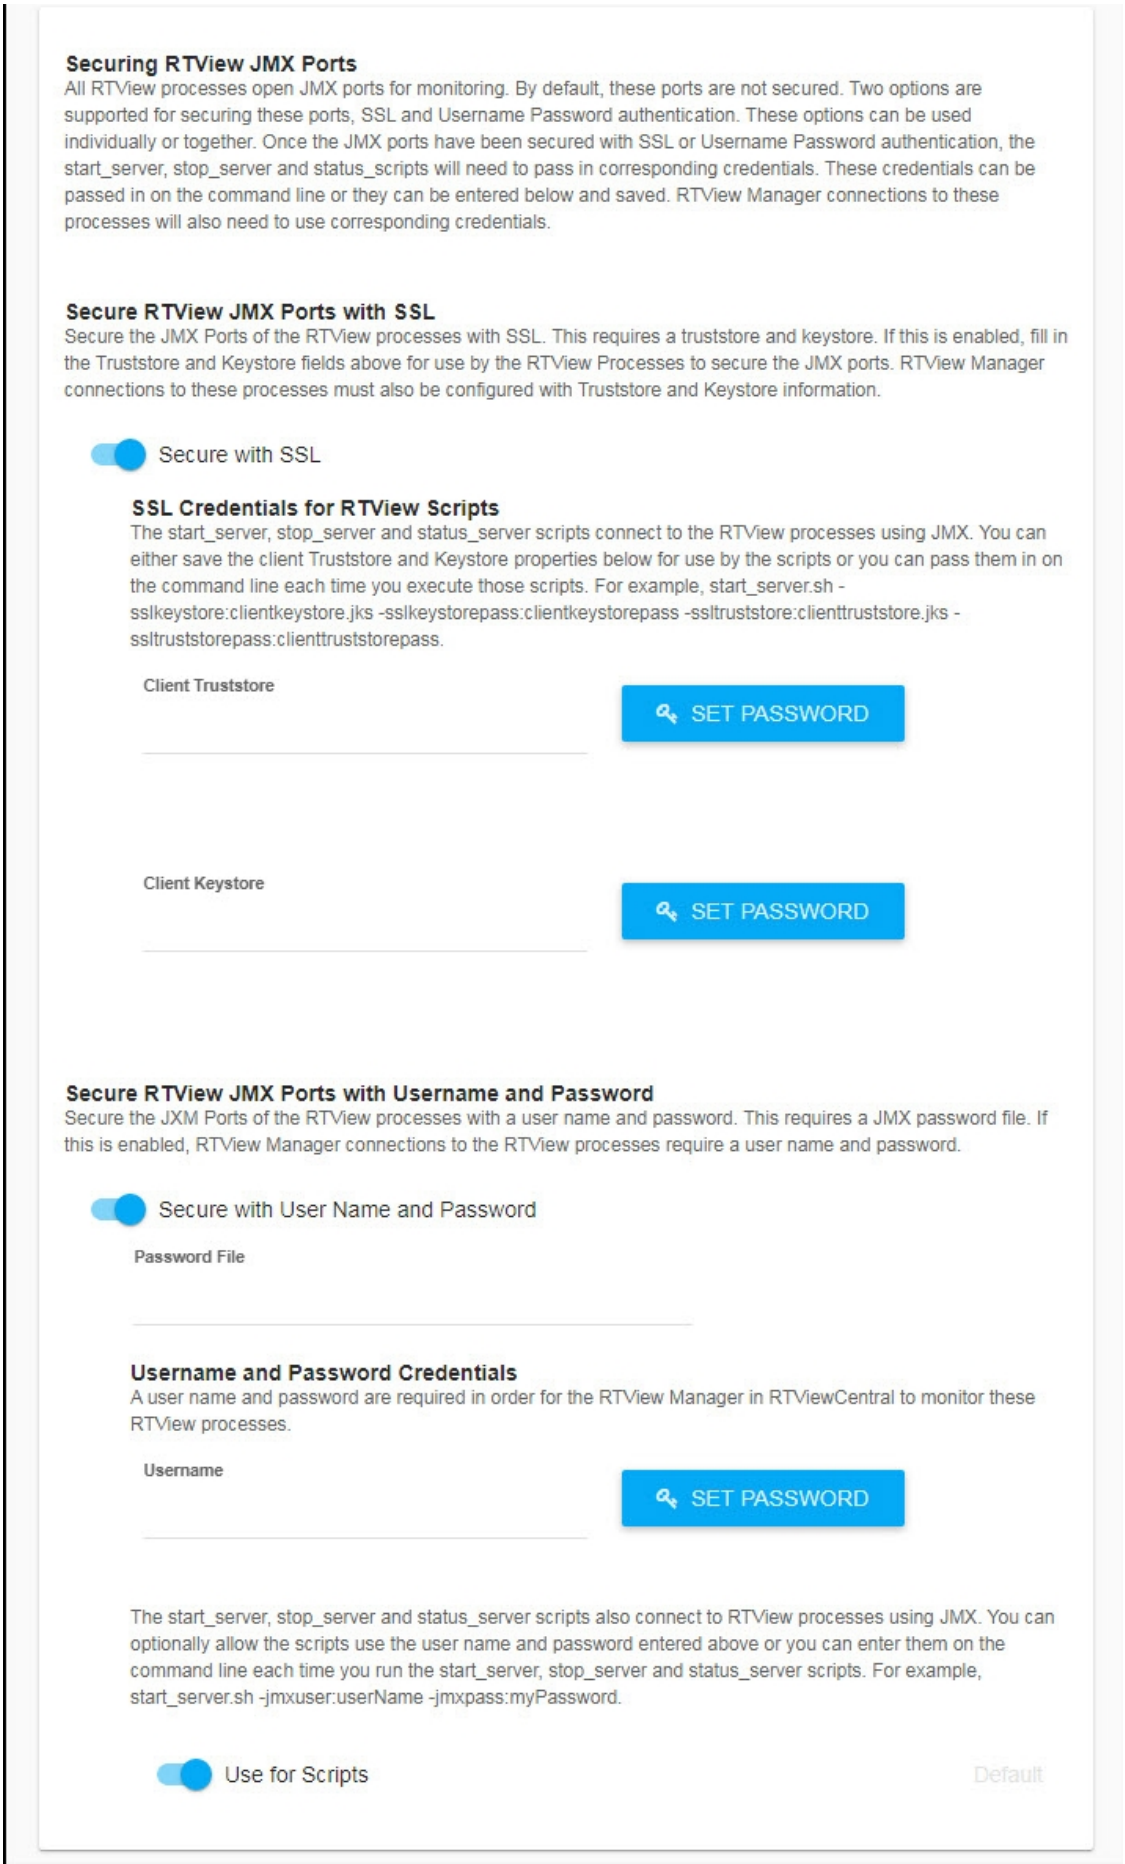
Securing RTView JMX Ports
All RTView processes open JMX ports for monitoring. By default, these ports are not secured. Two options are supported for securing these ports, SSL and Username Password authentication. These options can be used individually or together. Once the JMX ports have been secured with SSL or Username Password authentication, the start_server, stop_server and status_scripts will need to pass in corresponding credentials. These credentials can be passed in on the command line or they can be entered below and saved. RTView Manager connections to these processes will also need to use corresponding credentials.
Secure RTView JMX Ports with SSL
Secure the JMX Ports of the RTView processes with SSL. This requires a truststore and keystore. If this is enabled, fill in the Truststore and Keystore fields above for use by the RTView Processes to secure the JMX ports. RTView Manager connections to these processes must also be configured with Truststore and Keystore information.
Secure with SSL
SSL Credentials for RTView Scripts
The start_server, stop_server and status_server scripts connect to the RTView processes using JMX. You can either save the client Truststore and Keystore properties below for use by the scripts or you can pass them in on the command line each time you execute those scripts. For example, start_server.sh -sslkeystore:clientkeystore.jks -sslkeystorepass:clientkeystorepass -ssltruststore:clienttruststore.jks -ssltruststorepass:clienttruststorepass.
Client Truststore
SET PASSWORD
Client Keystore
SET PASSWORD
Secure RTView JMX Ports with Username and Password
Secure the JXM Ports of the RTView processes with a user name and password. This requires a JMX password file. If this is enabled, RTView Manager connections to the RTView processes require a user name and password.
Secure with User Name and Password
Password File
Username and Password Credentials
A user name and password are required in order for the RTView Manager in RTViewCentral to monitor these RTView processes.
Username
SET PASSWORD
The start_server, stop_server and status_server scripts also connect to RTView processes using JMX. You can optionally allow the scripts use the user name and password entered above or you can enter them on the command line each time you run the start_server, stop_server and status_server scripts. For example, start_server.sh -jmxuser:userName -jmxpass:myPassword.
Use for Scripts
Default

Secure Client and Receiver Ports with SSL

The Data Server opens a client port for use by other RTView processes and a receiver port to receive data from the Data Collector. By default, these port are not secured.

When **Secure Client Port with SSL** is enabled, the client port is SSL secured and all client data is encrypted using an anonymous cipher. However, no certificate is used to perform a SSL authentication and therefore the client and server do not verify each other's identities.

When **Secure RTView Receiver Port** is enabled, the receiver port is SSL secured and all receiver data is encrypted using an anonymous cipher. However, no certificate is used to perform a SSL authentication and therefore the client and server do not verify each other's identities.

Data Server

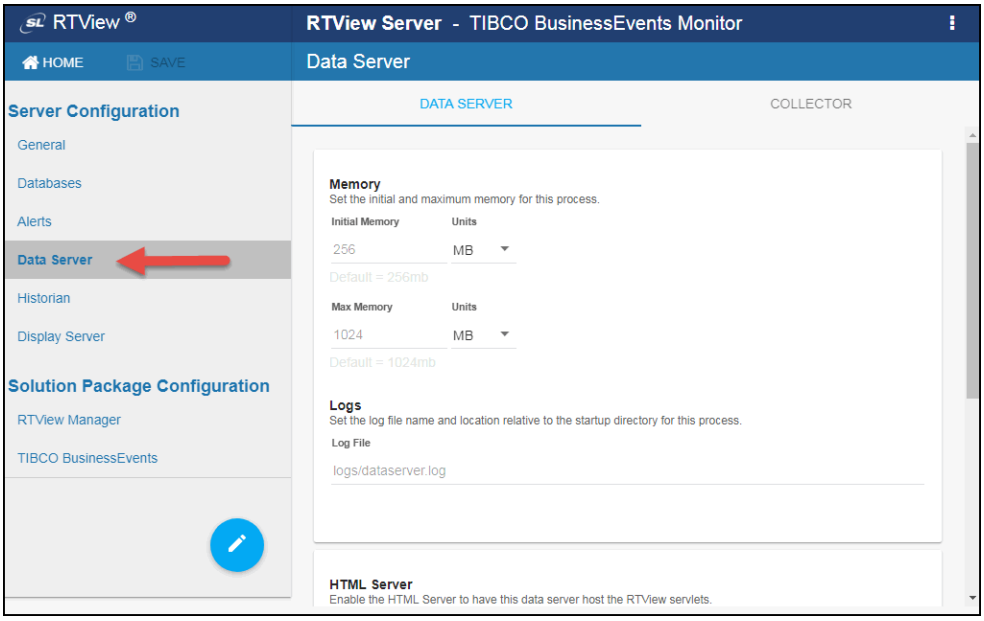

This section describes the Data Server Configuration settings. There are two tabs available:

- Data Server Tab
- Collector Tab

Data Server Tab

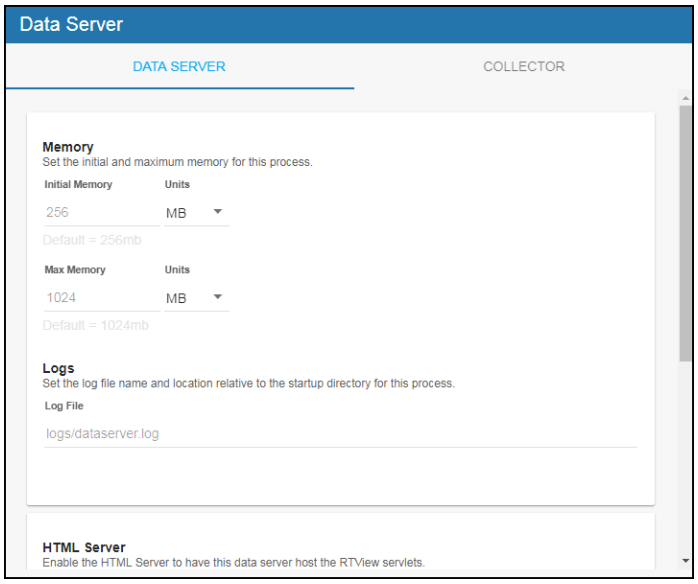

This tab contains the following:

Memory: Set the initial memory and maximum memory for the Data Server process. Specify a number followed by a unit. Units are k (kilobyte), m (megabyte), g (gigabyte). If no unit is used, the number is assumed to be bytes. **Note**: Use caution when you change the memory allocation. If the memory allocation is too small the server might crash during startup and if too large the server might eventually exceed the available CPU/memory and fail.

Initial Memory: The initial amount of memory to allocate for this process.

Max Memory: The maximum amount of memory to allocate for this process.

Logs

Log File: The log file name and location relative to the startup directory for this process. In the **Log File** field, use the following format: **<directory name>/<log file name>**.

For example, logs/dataserver.log.

HTML Server

HTML Server Enabled: Enable the Eclipse Jetty HTML Server in the Data Server. If enabled, Eclipse Jetty will host the RTView Servlets at **http://localhost:XX70**, where **XX** is the port prefix specified on the **Server Configuration** > **General** > **GENERAL** tab. **Note**: You cannot disable this option if the RTView Configuration Application is being hosted by Eclipse Jetty in the Data Server. All RTView Servlets hosted by Eclipse Jetty are automatically configured with the correct Data Server port at runtime. The following RTView Servlets are hosted in Eclipse Jetty:

rtview-tbemon-classic

rtview-tbemon-rtvadmin

rtvadmin

rtvdata

rtvquery

rtvagent

rtvpost

Collector Tab

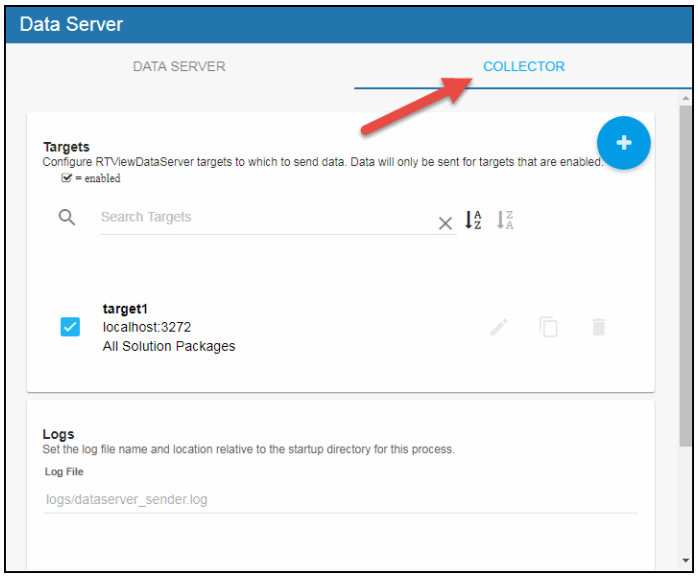

This tab should only be used when the data server is configured to be a sender. See Sender/Receiver: Distributing the Load of Data Collection for more information. This tab contains the following:

Targets: If you only have one target, you can edit the default target by clicking the "pencil" icon next to the **target1**. You can specify multiple targets by adding them one at a time. All fields on the **Add Target** dialog are required. Click the + icon to open the **Add Target** dialog, which has the following fields:

ID: A unique name for the target.

URL: Specify the URL for the receiver. The url can be **host:port** (for example, somehost:3372) or an **http url** for the rtvagent servlet on the receiver. For example, if you are using Tomcat, then you would use: **http://somehost:8068/tbemon-rtvagent**. If you are using Jetty, then you would use: **http://somehost:3270/rtvagent**.

Targets: Select the **All solution packages** option.

Enabled: Select this check box to enable the target.

Logs

Log File: The log file name and full path.

Identifier

Sender Name: A unique name for the sender data server, which is typically your machine's name.

Historian

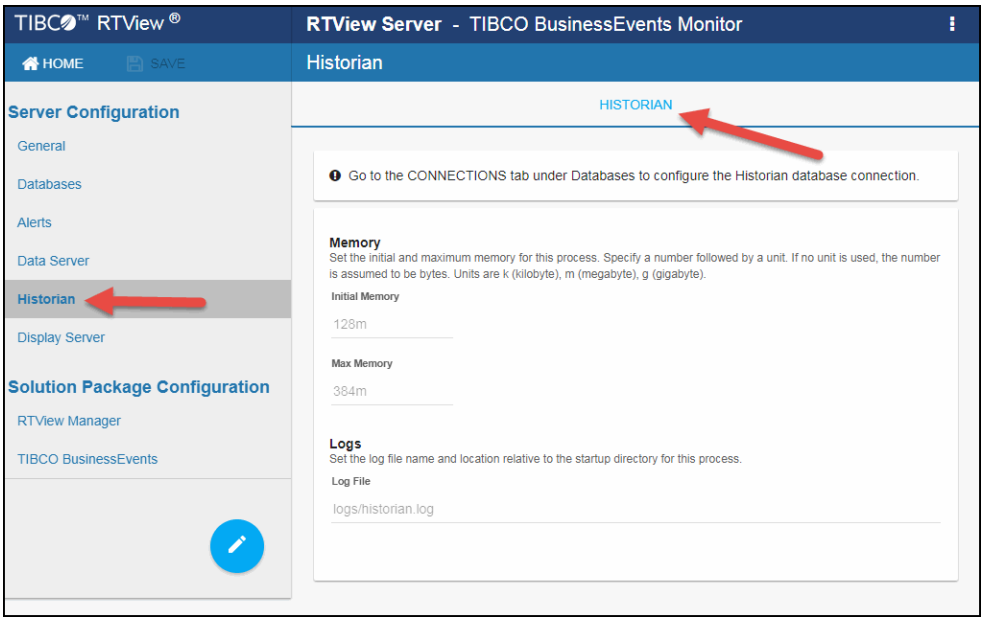

The **Historian** option consists of the **Historian** tab, which allows you to define the history properties. Historian database connections are set up using the Databases option. This option contains the following regions:

Memory: Set the initial memory and maximum memory for the Historian process. Specify a number followed by a unit. Units are k (kilobyte), m (megabyte), g (gigabyte). If no unit is used, the number is assumed to be bytes. **Note**: Use caution when you change the memory allocation. If the memory allocation is too small the server might crash during startup and if too large the server might eventually exceed the available CPU/memory and fail.

Initial Memory: The initial amount of memory to allocate for this process.

Max Memory: The maximum amount of memory to allocate for this process.

Logs

Log File: The log file name and location relative to the startup directory for this process. In the **Log File** field, use the following format: **<directory name>/<log file name>**.

For example, logs/historian.log.

Display Server

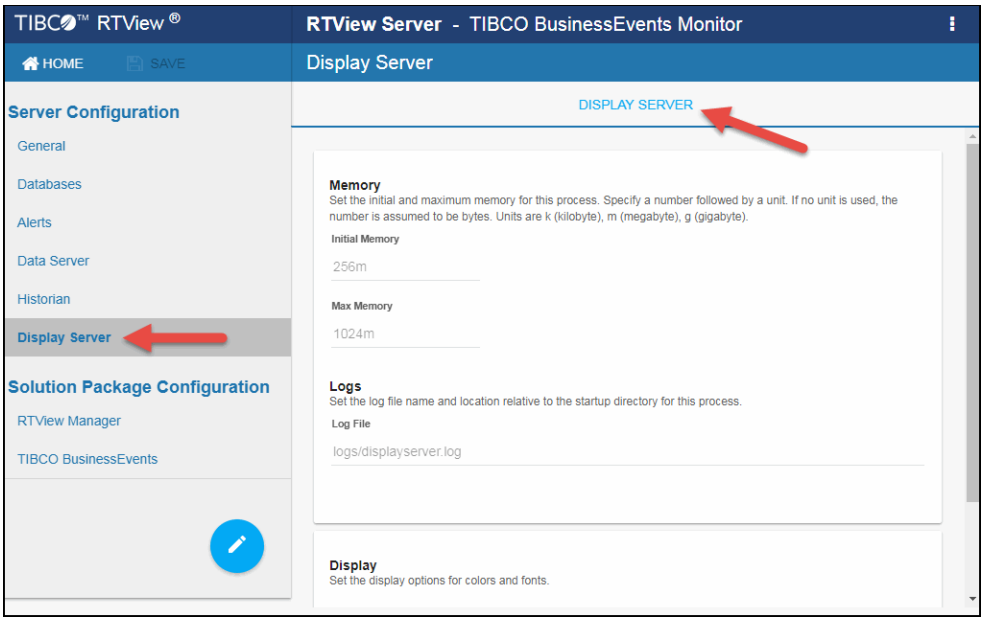

The **Display Server** option contains the **DISPLAY SERVER** tab, which contains the following regions:

Memory: Set the initial memory and maximum memory for the Display Server process. Specify a number followed by a unit. Units are k (kilobyte), m (megabyte), g (gigabyte). If no unit is used, the number is assumed to be bytes. **Note**: Use caution when you change the memory allocation. If the memory allocation is too small the server might crash during startup and if too large the server might eventually exceed the available CPU/memory and fail.

Initial Memory: The initial amount of memory to allocate for this process.

Max Memory: The maximum amount of memory to allocate for this process.

Logs

Log File: The log file name and location relative to the startup directory for this process. In the **Log File** field, use the following format: **<directory name>/<log file name>**.

For example, logs/displayserver.log.

Display

Use Light Stylesheet: By default, RTView displays appear with a black background. Select this option for all RTView displays to appear with a white background.

Enable Cross Platform Fonts: Selecting this option prevents LINUX users from seeing inconsistently aligned labels in displays. This option should only be applied to Display Servers on Linux AND only if the text size and alignment issue is observed in the Thin Client. Otherwise, it can cause unnecessary overhead or unwanted changes to the appearance of text in RTView displays.

Solution Package Configuration View

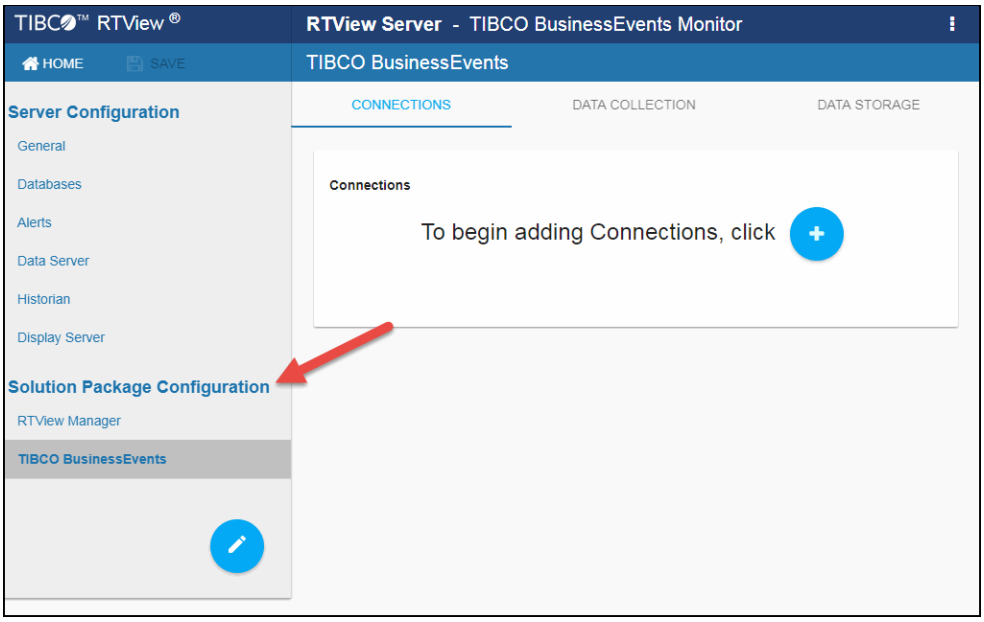

The **Solution Package Configuration** View provides options that allow you to modify the default settings for the project, define the classpaths and connections for the Monitor, and define the data collection and data storage properties for the Monitor. Descriptions for all of the properties for these options, as they pertain to the Monitor, are explained in detail in the Configuration chapter. You can also view the basic steps to get the Monitor up and running in the Quick Start chapter.

APPENDIX D Security Configuration

This section provides details for securing a direct connection RTView deployment. This section contains:

- Introduction
- Data Server
- Display Server (thin client)
- HTML UI
- Data Collectors
- Configuration Application
- Configuration Files
- Historian
- Database
- Application Servers
- Monitored Components
- Security Summary

Introduction

The following diagram shows how data flows through the RTView deployment. The Data Server connects to the Monitored Components to collect metric data which it stores in local caches. The Data Server uses the collected data to generate alerts based on enabled threshold settings in the ALERDEFS database. If the user has (optionally) defined alert notifications, the Data Server also executes them.

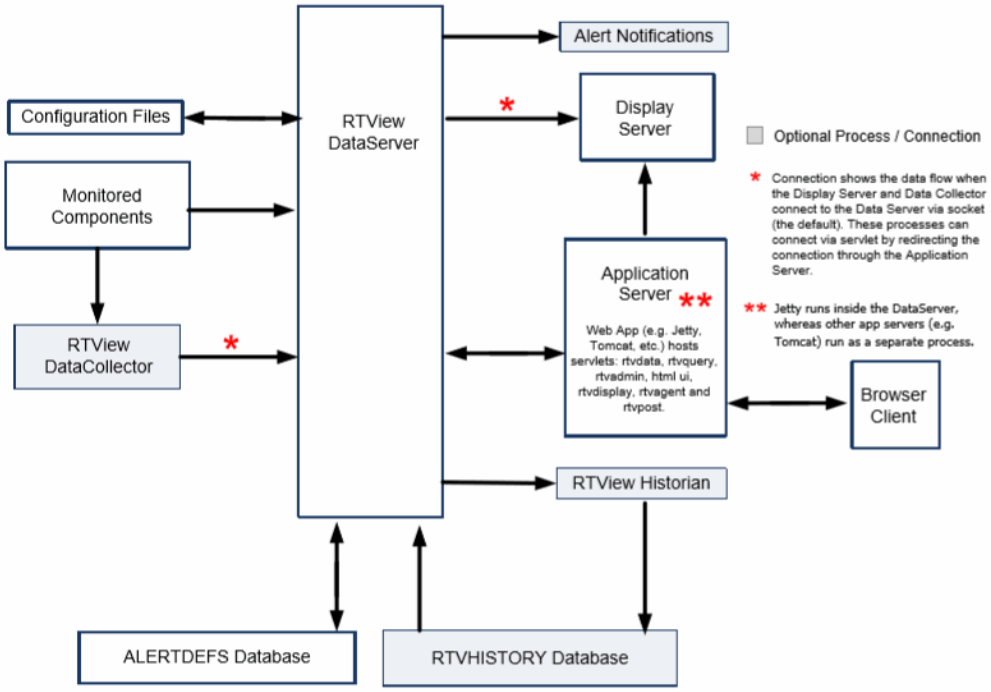

In cases where the data collection needs to be distributed, one or more Data Collectors can be deployed to connect to the Monitored Components and forward the collected data to the Data Server.

The HTML UI and Display Server (thin client) are browser-based user interfaces that show metric and alert data from the Data Server and also allow the user to enable, disable and set thresholds on alerts.

The Historian is an optional process that stores historical metric and alert data to the RTVHISTORY database. When the Historian is enabled, the Data Server queries historical data from the RTVHISTORY on startup to populate in-memory history and also any time the Display Server or HTML UI request history data that is older than the data in the in-memory history.

The Configuration Application is a browser based application for configuring the RTView processes. It connects to the Data Server to read and write Configuration Files.

The next sections provide a more detailed description of each process.

Port Settings

This document describes port settings for the following TIBCO applications, where the **XX** prefix is replaced with the following:

- For TIBCO EMS, replace **XX** with **31**
- For TIBCO BusinessWorks5, replace **XX** with **33**
- For TIBCO BusinessWorks6, replace **XX** with **45**
- For TIBCO BusinessEvents, replace **XX** with **32**

Data Server

The Data Server gathers and caches the data from the applications being monitored and also hosts the alerts for that data. Because the Data Server can exist behind firewalls, it simplifies and strengthens the secured delivery of information to clients beyond the firewall. The Data Server serves metric and alert data to the Display Server and Historian via socket on port **XX78** and receives data via socket from the optional Data Collector on port **XX72**. It also serves metrics and alert data to the HTML UI via the rtvquery servlet which connects via socket on port **XX78**.

The Historian and Display Server run in the same directory as the Data Server, while the optional Data Collector(s) typically run in a different directory or a different system. By default, socket connections to the Data Server are unsecured. The Data Server supports secure socket connections (SSL) with or without certificates. It also supports client whitelist and blacklist. Secure socket and client whitelist/blacklist configuration are described in the *RTView Core User's Guide* under Deployment/Data Server/Security.

The Display Server can optionally be configured to connect to the Data Server via the rtvdata servlet instead of the socket. In this case, the rtvdata servlet connects to the Data Server via socket on port **XX78**. While the rtvdata servlet cannot be configured for authentication, Tomcat access filters can be used to restrict access. The rtvdata servlet will connect via secure socket if the Data Server is configured for SSL sockets.

The HTML UI connects to the Data Server via the rtvquery servlet. See HTML UI in this document for information on how to enable authentication in the HTML IU and rtvquery servlets. The rtvquery servlet will connect via secure socket if the Data Server is configured for SSL sockets.

The Data Collector can optionally be configured to send data to the Data Server via the rtvagent servlet instead of the socket. In this case, the rtvagent servlet connects to the Data Server via socket on port **XX72**. While the rtvagent servlet cannot be configured for authentication, Tomcat access filters can be used to restrict access. The rtvagent servlet will connect via secure socket if the Data Server is configured for SSL sockets.

The Configuration Application connects to the Data Server via the rtvadmin servlet to read and write properties files. The rtvadmin servlet connects to the Data Server via socket on port **XX78**. See Configuration Application in this document for information about servlet authentication. The rtvadmin servlet will connect via secure socket if the Data Server is configured for SSL sockets.

If the Historian is enabled, the Data Server connects to the RTVHISTORY database on startup to read initial cache history data and if the thin client or HTML UI request history data older than the in memory cache history. It also connects to the ALERTDEFS database to query and set alert thresholds. See Database in this document for more information.

The Data Server optionally executes alert notifications based on user settings. Since the notification actions are user defined, security must be determined by the user.

The Data Server opens a JMX port on **XX68** to enable monitoring. By default, the JMX port is not secured. See RTView Processes under Monitored Components for information on securing this connection.

By default, the Data Server runs a Jetty process which hosts all of the RTView servlets and accepts HTTP client requests on port **XX70**. You can optionally configure Jetty to use HTTPS instead of HTTP.

Also see Port Settings.

Display Server (thin client)

The Classic user interface deployment, the thin client, is implemented using the Display Server. The Display Server consists of two parts, the Display Server application and the Display Servlet (**rtview-<sp>-classic.war**, where **<sp>** is replaced with **emsmon**, **bwmon**, **bw6mon** or **tbemon**). The Display Server application is generally installed on a dedicated platform. It loads displays from the file system and queries data from the Data Server that it passes on to the Display Servlet via a socket. The Display Servlet runs on an application server (like Tomcat or Jetty). Browser clients connect to the Display Servlet using HTTP or HTTPS (depending on the Application Server configuration). This process opens a JMX port on **XX79** to enable monitoring. By default, the JMX port is not secured. See RTView Processes under Monitored Components for information on securing this connection.

Also see Port Settings.

For the thin client, we support our **users**/**role.xml** login (which can be customized to integrate with LDAP) as well as SSO using BASIC or DIGEST HTTP authentication. This is described in the *RTView Core User's Guide* under Role-based Security and also under Deployment/Browser Deployment/Display Server.

> **Note:** The external libraries used for PDF export capability, iText and iTextAsian, can trigger security warnings due to a High Severity CVE (http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2017-9096) which "might allow remote attackers to conduct XML external entity attacks via a crafted PDF".
>
> These libraries do not pose a genuine security risk. The only PDFs that are handled by iText in RTView are those generated from customer displays, and therefor there is no mechanism to exploit the vulnerability with a "crafted PDF".
>
> Users who do not need to use the PDF export or reporting capability, and who would like to remove any library that raises alarms in a security scan, can remove gmsjextpdf.jar from the lib directory.

HTML UI

The new user interface is implemented in HTML and is deployed as a servlet, **rtview-<sp>mon** (where **<sp>** is replaced with **emsmon**, **bwmon**, **bw6mon** or **tbemon**), which is configured by default to use BASIC HTTP authentication. Browser clients connect via HTTP or HTTPS depending on the Application Server configuration. It should be used with HTTPS since BASIC authentication does not encrypt user credentials. The HTML UI connects to the Data Server via the rtvquery servlet. See Data Server for information on securing the connection between the rtvquery servlet and the Data Server. By default, the rtvquery servlet is not configured for authentication, but can be modified to require BASIC HTTP authentication as follows (this should be used with HTTPS since BASIC authentication does not encrypt user credentials):

1. Extract the web.xml file from the rtvquery servlet as follows:

   jar -xf rtview-<sp>-rtvquery.war WEB-INF/web.xml

   where <sp> is replaced with emsmon, bwmon, bw6mon or tbemon
2. Open **WEB-INF/web.xml** in a text editor and uncomment the security section.
3. Pack the modified **web.xml** file back into the rtvquery servlet as follows:

   jar -uf rtview-<sp>-rtvquery.war WEB-INF/web.xml

   where <sp> is replaced with emsmon, bwmon, bw6mon or tbemon

After you enable BASIC HTTP authentication in the rtvquery servlet, you will also need to modify the HTML UI to pass in credentials:

1. Extract the **setup.js** file from r**tview-<sp>.war** as follows:

   jar -xf rtview-<sp>.war setup.js

   where <sp> is replaced with emsmon, bwmon, bw6mon or tbemon
2. Open **setup.js** in a text editor and remove the **//** from the beginning of the **authValueC** line: **//authValueC: 'Basic ' + btoa('rtvuser:rtvuser')**
3. Pack the modified **setup.js** file back into the HTML UI servlet as follows:

   jar -uf rtview-<sp>.war setup.js

   where **<sp>** is replaced with **emsmon**, **bwmon**, **bw6mon** or **tbemon**

Data Collectors

This process is optional and is used to distribute connections to Monitored Components Data Collectors intead of having the Data Server connect to each component to be monitored directly. This process collects data from Monitored Components and forwards it to the Data Server via socket or the rtvagent servlet. See Data Server for information about securing the connection between the Data Collector and Data Server. This process does not keep history or process alerts - those are handed in the Data Server. While the Data Collector typically does not have data clients, it accepts data requests via socket on port **XX76** which can be secured as described in the Data Server section. It runs Jetty on port **XX70** and also opens JMX on port **XX66** for monitoring. By default, the JMX port is not secured. See RTView Processes under Monitored Components for information on securing this connection.

Also see Port Settings.

Configuration Application

The Configuration Application connects to the Data Server via the rtvadmin servlet which is configured with BASIC HTTP authentication. It should be run on HTTPS since Basic Authentication does not encrypt user credentials. Passwords saved by the configuration application are scrambled except in the case where they are added in the **CUSTOM PROPERTIES** section. See Data Server for information about securing the connection between the Configuration Application and Data Server.

Configuration Files

Configuration (**.properties** and **.properties.json**) files are stored on the file system and read by all RTView processes to control configuration. Additionally, the Configuration Application writes these files, scrambling all connection and database passwords. Passwords entered in the **CUSTOM PROPERTIES** tab are not scrambled.

Historian

The Historian connects to the Data Server via socket and saves cache history to a database via JDBC. This process is optional and the user can configure which data will be saved. See Data Server for information about securing the connection between the Historian and Data Server. See Database for information about the connection between the Historian and the database.

This process opens JMX port **XX67** for monitoring. By default, the JMX port is not secured. See RTView Processes under Monitored Components for information on securing this connection.

Also see Port Settings.

Database

The ALERTDEFS database stores alert threshold information and optionally alert persistence information. The Data Server connects to the ALERTDEFS database to query thresholds and also to set thresholds when the user interacts with the **Alert Administration** page in the user interface. The RTVHISTORY database stores cache data (if the Historian is enabled). The Historian connects to the RTVHISTORY database to insert cache history data and to perform data compaction. The Data Server connects to the RTVHISTORY database on startup to load initial history into the caches and also when the user interface asks for history data older than what is contained in the in-memory history caches.

By default, the Data Server and Historian connect to the HSQLDB database that is included with RTView using an unsecured JDBC connection. See the HSQLDB documentation for information on configuring it for secure JDBC connections. Alternately, you can use your own database and secure the JDBC connection according to the documentation for that database.

Application Servers

By default, the Data Server runs a Jetty process which hosts all of the RTView servlets and accepts HTTP client requests on port **XX70**. You can optionally configure Jetty to use HTTPS instead of HTTP. This will require you to provide a certificate for your domain.

Also see Port Settings.

When you have a certificate, do the following in the Configuration Application in the **Data Server** tab:

1. Turn on the **Use HTTPS** toggle.
2. Set the **Keystore File** to the keystore file name (including the path) that contains the certificate for your domain.
3. Optionally enter the **Keystore Password** and **Key Manager Password** if they are required for your keystore.
4. **Save** your configuration and restart the data server.

The Configuration Application and HTML UI use BASIC HTTP authentication and require the following roles which are preconfigured. You can modify the user names and passwords (but not the roles) in **RTVAPM_HOME/common/lib/ext/jetty/rtvadmin-users.xml**:

- rtvadmin
- rtvuser
- rtvalertmgr

Jetty does not limit the number of failed login attempts which leaves it open to brute force attacks. If this is a concern, you should deploy with Tomcat or another Application Server.

You can optionally use Tomcat or another Application Server in addtion to or instead of the Jetty process that comes with RTView. To deploy your servlets to your application server, go into the **RTVAPM__HOME/<sp>/projects/sample** directory (where **<sp>** is replaced with **emsmon**, **bwmon**, **bw6mon** or **tbemon)** and run **update_wars.bat** or **update_wars.sh**. Copy all of the generated war files to the **webapps** directory in your application server.

Tomcat and most other Application Servers can be configured for HTTPS. This will require you to provide a certificate for your domain. Follow the application server instructions to enable HTTPS.

Additionally, Tomcat access filters can be configured to restrict access according to the remote client host or address. Tomcat also has a feature named LockOut Realm to protect against brute force login attacks. After 5 successive login attempts for a given username with invalid password, then all logins for that username are rejected for the next 5 minutes. The LockOut Realm parameters are configurable. See Apache Tomcat documentation for more information.

You will need to add the following roles to your Application Server for use with the Configuration Application and HTML UI authentication. For Tomcat, users and roles are defined in **conf\tomcat-users.xml**:

- rtvadmin
- rtvuser
- rtvalertmgr

You can optionally disable Jetty in the Data Server when using Tomcat or another Application Server. To disable Jetty, you must access the Configuration Application from Tomcat or another Application Server. In the Configuration Application, go to the **Data Server** tab and do the following:

- Turn off the **HTML Server Enabled** toggle.
- **Save** your configuration and restart.

Monitored Components

Monitored Components are the processes that the Data Server and Data Collector connect to in order to request metric data. Some examples of Monitored Components are EMS Servers, Oracle Databases and RTView Processes. Connections to Monitored Components are made through application-specific APIs, so the options for securing these connections differ based on the Monitored Component.

This section contains:

- TIBCO BusinessEvents
- TIBCO BusinessWorks 5
- TIBCO BusinessWorks 6
- TIBCO EMS Server
- TIBCO Hawk
- RTView Manager
- RTView Processes

TIBCO BusinessEvents

The Data Server connects to BusinessEvents using JMX. BusinessEvents does not support secure JMX connections.

TIBCO BusinessWorks 5

The Data Server connects to TIBCO BusinessWorks 5 using TIBCO Hawk. See TIBCO Hawk for information about securing those connections. Additional server metrics can optionally be collected via JMX using the RTView Manager . See the TIBCO BusinessWorks 5 documentation for enabling JMX and securing it in your TIBCO BusinessWorks engine. See RTView Manager in

this document for information on making secure connections to JMX. For BWSE engines, RTView collects AMX Node data via JMS messages which are hosted on an EMS Server. The EMS Server can be configured to require a user name and password which the user enters in the Configuration Application when you define the Connection to that Server. Additionally, the EMS Server can be configured to use SSL. In this case, the user must implement a subclass of the GmsRtViewJmsDsSSLHandler to return a Map of the required SSL parameters per connection. This is described in the *RTView Core User's Guide* under **RTView Data Sources/JMS Data Source/Application Options - JMS/JMS Connections Tab/JMS SSL Parameters**.

TIBCO BusinessWorks 6

The Data Server connects to BusinessWorks 6 either using TIBCO Hawk or via the OSGI plugin. See the TIBCO Hawk section for information about securing TIBCO Hawk connections. When using the OSGI plugin, data is sent via socket to the Data Server on port **XX72**. By default, this socket is not secure, but the data will be sent via secure socket if the Data Server is configured for SSL sockets.

Also see Port Settings.

TIBCO EMS Server

The Data Server connects to EMS Servers using TIBCO's TibjmsAdmin API. The EMS Server can be configured to require a user name and password which the user enters in the Configuration Application/**EMS Server Connection** dialog when you define the connection to that EMS Server. Additionally, the EMS Server can be configured to use SSL. In this case, the user must implement a subclass of the GmsRtViewTibJmsSSLHandler to return a Map of the required SSL parameters per connection. This is described in the *RTView Core User's Guide* under **RTView Data Sources/TIBCO EMS Administration Data Source/Application Options - TIBCO EMS/TIBCO EMS Servers Tab/TIBCO EMS Administration SSL Parameters**.

TIBCO Hawk

TIBCO Hawk is used to gather metrics for both BusinessWorks 5 (required) and BusinessWorks 6. The Data Server connects to TIBCO Hawk via TIBCO's TIBHawkConsole API. The TIBCO Hawk installation may either be configured to run on a Rendezvous transport (rvd) or an EMS Transport. In the case of Rendezvous transports, no secure connection options are supported. In the case of EMS transports, the TIBCO Hawk agent can be configured to require a user name and password which the user enters in the Configuration Application **TIBCO Hawk Connection** dialog when adding a connection to that agent. The EMS transport can also be configured for SSL. In this case, the user must implement a subclass of the GmsRtViewHawkCustomSSLHandlerto return a Map of the required SSL parameters per connection. This is described in the *RTView Core User's Guide* under **RTView Data Sources/TIBCO Hawk Data Source/Application Options - TIBCO Hawk/TIBCO Hawk SSL Parameters**.

RTView Manager

The RTView Manager connects to processes via JMX. A process that opens a JMX port can be configured to require a user name and password which the user enters in the Configuration Application **RTView Manager Connection** dialog when you define the connection to that process. A process that opens a JMX port can also be configured to require SSL. To connect to

SSL JMX, in the **CUSTOM PROPERTIES** tab, add the following properties replacing the values with the appropriate values for your SSL configuration:

1. Set key store:
   a. **Property Name**: sl.rtview.jvm
   b. **Propety Value**: -Djavax.net.ssl.keyStore=client_keystore.jks
2. Set keystore password:
   a. **Property Name**: sl.rtview.jvm
   b. **Property Value**: -Djavax.net.ssl.keyStorePassword=mypassword
3. Set truststore
   a. **Property Name**: sl.rtview.jvm
   b. **Property Value**: sl.rtview.jvm=-Djavax.net.ssl.trustStore=client_truststore.jks
4. Set truststore password
   a. **Property Name**: sl.rtview.jvm
   b. **Property Value**: sl.rtview.jvm=-Djavax.net.ssl.trustStorePassword=mypassword
5. **Save** your changes and **Restart Servers**.

RTView Processes

The Data Server, Data Collector, Historian and Display Server all open JMX ports for monitoring. By default, these JMX ports are unsecured, but they can be secured either by user name and password or by SSL. See RTView Manager in this document for instructions on connecting to secure JMX. Note that the **start**/**stop**/**status_rtv** scripts use JMX to communicate with the RTView processes. If you secure the JMX ports on the RTView processes, you will need to pass additional arguments into the **start**/**stop**/**status_rtv** scripts as described in the *RTView Enterprise Monitor Configuration Guide*.

To secure with user name and password:

1. Setup a password file in **JRE_HOME/lib/management** using the **jmxremote.password.template** file as a template as described in the Oracle docs (for example, you might refer to Using File-Based Password Authentication here: https://-docs.oracle.com/javase/8/docs/technotes/guides/management/agent.html).
2. Start the Data Server for the project you want to secure and open the Configuration Application. In the **CUSTOM PROPERTIES** tab, add the following properties:
   - Disable ssl

   Property Name: sl.rtview.jvm

   Property Value: -Dcom.sun.management.jmxremote.ssl=false
   - Enable authentication:

   Property Name: sl.rtview.jvm

   Property Value: -Dcom.sun.management.jmxremote.authenticate=true
   - Set password file:

   Property Name: sl.rtview.jvm

   Property Value: -Dcom.sun.management.password.file=jmxremote.properties (or the name and path to a different password file)
3. **Save** your changes and **Restart Servers**. The JMX Port on all RTView Servers for that project will now require a user name and password.

To secure with SSL:

If you do not already have a key pair and certificate setup on your server, follow the instructions in the Oracle docs (for example, you might refer to Using SSL/To Setup SSL, Step 1 here: https://docs.oracle.com/javase/8/docs/technotes/guides/management/agent.html).

1. Start the Data Server for the project you want to secure and open the Configuration Application. In the **CUSTOM PROPERTIES** tab, add the following properties replaceing the values with the appropriate values for your SSL configuration:
   - Enable ssl

   Property Name: sl.rtview.jvm

   Property Value: -Dcom.sun.management.jmxremote.ssl=true
   - Set key store:

   Property Name: sl.rtview.jvm

   Propety Value: -Djavax.net.ssl.keyStore=server_keystore.jks
   - Set keystore password:

   Property Name: sl.rtview.jvm

   Property Value: -Djavax.net.ssl.keyStorePassword=mypassword
   - Set truststore

   Property Name: sl.rtview.jvm

   Property Value: sl.rtview.jvm=-Djavax.net.ssl.trustStore=server_truststore.jks
   - Set truststore password

   Property Name: sl.rtview.jvm

   Property Value: sl.rtview.jvm=-Djavax.net.ssl.trustStorePassword=mypassword
2. **Save** your changes and **Restart Servers**. The JMX Port on all RTView Servers for that project will now require SSL.

Security Summary

Security options per RTView process are included in the section for each component above. This section provides a summary of security options for the entire deployment organized by priority.

This section contains:

- Secure Installation Location - High Priority
- Login and Servlet Authentication - High Priority
- Application Server Security - High Priority
- Secure Connections between RTView Processes - Medium-to-Low Priority*
- Secure Connections to Monitored Components - Medium-to-Low Priority*
- Secure Connections to Monitored Components - Medium-to-Low Priority*

Secure Installation Location - High Priority

The RTView installation and Application Server should be run in a secure location to ensure displays and configuration files are secure and access-restricted.

Login and Servlet Authentication - High Priority

- **HTML UI** - By default, the HTML UI is configured with BASIC HTTP authentication which should use HTTPS since BASIC authentication does not encrypt user credentials. The HTML UI connects to the Data Server via the rtvquery servlet. The rtvquery servlet does not have authentication enabled by default. See the HTML UI section in this document for information on enabling authentication in the rtvquery servlet.
- **Display Server** - By default, the Display Server authentication is disabled. See the Display Server (thin client) section in this document for information on enabling authentication for the Display Server.
- **Configuration Application** - By default, the Configuration Application is configured with BASIC HTTP authentication which should use HTTPS since BASIC authentication does not encrypt user credentials.

Application Server Security - High Priority

It is highly recommended that you configure your Application Server to use HTTPS as described in the Application Servers section of this document. The RTView servlets that support HTTP authentication all use BASIC authentication which does not encrypt user credentials.

It is highly recommended that you change the user credentials in your Application Server for the rtvadmin, rtvuser and rtvalertmgr roles since the default credentials are documented and publicly available.

Secure Connections between RTView Processes - Medium-to-Low Priority*

The Historian, Data Server, Data Collector, rtvquery servlet, rtvdata servlet, rtvadmin servlet and rtvagent servlet all connect to the Data Server via socket which is unsecured by default. The Data Server supports secure socket connections (SSL) with or without certificates. It also supports client whitelist and blacklist. Secure socket and client whitelist/blacklist configuration are described in the RTView Core User's Guide under **Deployment/Data Server/Security**.

Secure Connections to Monitored Components - Medium-to-Low Priority*

The Data Server uses component specific API's to connect to Monitored Components. See the Monitored Components section in this document for information on how to secure these connections.

Secure Connections to Databases - Medium-to-Low Priority*

The Data Server and Historian both create database connections using JDBC. See the Database section in this document for information on securing JDBC connections to your database.

*If Secured Installation Location has been met, these are lower priority.