RTView® Monitor for Solace® User's Guide

Version 4.2

RTView Enterprise Monitor®

Copyright © 2013-2018. Sherrill-Lubinski Corporation. All rights reserved.

RTView®

Copyright © 1998-2018. Sherrill-Lubinski Corporation. All rights reserved.

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

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

SL, SL-GMS, GMS, RTView, RTView Core, RTView Enterprise Monitor, SL Corporation, and the SL logo are trademarks or registered trademarks of Sherrill-Lubinski Corporation in the United States and other countries. Copyright © 1998-2017. Sherrill-Lubinski Corporation. All rights reserved.

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

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

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

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

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

Contents

Contents iii

Preface 1
About This Guide 1
Document Conventions 1
Additional Resources 1
Release Notes 2
Documentation and Support Knowledge Base 2
Contacting SL 2
Internet 2
Technical Support 2

Chapter 1 - Quick Start 3
Prerequisites for Solace On-premise Installations 3
Prerequisites for Solace AMI 3
On-premise Initial Steps 4
Add Message Router Connections 4

Chapter 2 - Introduction to the Monitor 9
Overview 9
RTView Monitor for Solace On-Premise Version 10
RTView Monitor for Solace AMI Version 10
System Requirements 10
Installation 10
File Extraction Considerations 11

Chapter 3 - Configuration 13
Overview 13
Initialize a Command Prompt or Terminal Window 13
Configure Data Collection 14
Open the RTView Monitor for Solace Project 14
Define Solace Message Router and Syslog Connections 15
Modify Default Settings for Storing Historical Data 20
Change Port Assignments 25
Configure the Database 27

Configure Alert Notifications 31
Configuring Alert Notifications using the RTView Configuration Application 31
Configuring Monitor Alert Notification Actions 35
Using a Batch File or Shell Script 35
Using the Java Command Handler 37
Customizing the Java Command Handler 37
Java Command Handler Substitutions 37
Troubleshooting 38
Log Files for Solace 38
JAVA_HOME 38
Permissions 38
Network/DNS 38
Data Not Received from Data Server 39
Stop the Monitor 39
Start the Monitor 39

Chapter 4 - Additional Configurations 41
Obtain SEMP Version 41
Create Instance from RTView Monitor for Solace 42

Chapter 5 - Using the Monitor 47
Overview 48
Heatmaps 48
Mouse-over 50
Tables 50
Multiple Column Sorting 51
Column Visibility 51
Column Filtering 51
Column Locking 53
Column Reordering 53
Saving Settings 54
Row Paging 54
Trend Graphs 55
Time Settings 55
Mouse-over 55
Log Scale 56
GUI Icons and Buttons 56
RTView Monitor for Solace Views/Displays 58
Routers 59
Routers - Overview 59
Routers - Heatmap 61
Routers - Table 63
Routers - Summary 70
Routers - Sensors 72
Routers - Provisioning 74

Routers - Interface 75
Routers - Message Spool 77
Neighbors 80
Neighbors - Table 80
Neighbors - Diagram 81
Neighbors - Summary 82
VPNs 84
VPNs - Heatmap 84
VPNs - Table 88
VPNs - Summary 92
Clients 94
Clients - Table 94
Client - Summary 99
Bridges 102
Bridges - Table 102
Bridges - Diagram 106
Bridge - Summary 107
Endpoints 110
Endpoints - Table 110
Endpoint - Summary 112
Capacity 115
Capacity - Table 115
Capacity - Summary 117
Capacity - Trends 119
Alerts Table 120
..... 124
Administration 124
Alert Administration 124
Setting Override Alerts 128
Alert Administration Audit 129
RTView Cache Tables 131
RTView Agent Admin 133
RTView Manager for Solace Displays 133
Syslog 134
All Syslog Events Table 134
Alert Views 136
Alert Detail Table 136
Administration 139
Alert Administration 140
Tabular Alert Administration 141
Setting Override Alerts 143
Alert Administration Audit 144
RTView Agent Metrics Administration 147
RTView Cache Tables 148
RTView Agent Administration 149
About 150

RTView Manager Views/Displays 151
JVM Process Views 151
All JVMs Heatmap 152
All JVMs Table 154
JVM Summary 156
JVM System Properties 160
JVM Memory Pool Trends 161
JVM GC Trends 165
RTView Servers 167
Data Servers 167
Display Servers 170
Historian Servers 171
Version Info 173
Tomcat Servers 175
All Tomcat Servers 175
Tomcat Server Summary 178
All Applications Heatmap 180
Single Application Summary 182
MySQL Database 185
All Servers Heatmap 185
All Servers Table 188
Server Summary 190
Servers Properties 192
Servers Operations 193
User Tables 195
Docker Engines 196
Engines Heatmap 196
Engines Table 199
Engine Summary 201
Containers Heatmap 203
Containers Table 206
Container Summary 208
Hosts 211
All Hosts Heatmap 211
All Hosts Table 213
All Hosts Grid 215
All Processes Table 217
All Network Table 219
All Storage Table 221
Host Summary 223
Alert Views 225
Alert Detail Table 225
Administration 230
Alert Administration 230
Alert Administration Audit 237
Metrics Administration 238
RTView Cache Tables 240

RTView Agent Admin 242

Appendix A - RTView Configuration Application 245

Open the RTView Configuration Application 245

HOME Page 246

Apply Changes 246

Settings for RTView Central Servers 247

Central Server Configuration>General 247

Central Server Configuration>Data Servers>CONNECTIONS 250

Central Server Configuration>Central Config Server 251

Central Server Configuration>Central Alert Server 255

Central Server Configuration>Central Alert Historian 259

Central Server Configuration>Central Display Server 260

Settings for Solution Package Servers 261

Solution Package Server Configuration>General 261

Solution Package Server Configuration>Data Server 264

Solution Package Server Configuration>Data Server>DATA SERVER Tab 265

Solution Package Server Configuration>Data Server>SENDER Tab 265

Solution Package Server Configuration>Historian 266

Solution Package Server->Solution Package Configuration 267

Appendix B - Monitor Scripts 269

Scripts 269

rtvservers.dat 277

Appendix C - Alert Definitions 281

Appendix D - Third Party Notice Requirements 287

Appendix E - Limitations 305

iPad Safari Limitations 305

Preface

Welcome to the *RTView® Monitor for Solace® User's Guide*.

Read this preface for an overview of the information provided in this guide and the documentation conventions used throughout, additional reading, and contact information. This preface includes the following sections:

- "About This Guide" on page 1
- "Additional Resources" on page 1
- "Contacting SL" on page 2

About This Guide

The *RTView® Monitor for Solace® User's Guide* describes how to install, configure and use the Monitor.

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>**

Additional Resources

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

- "Release Notes" on page 2
- "Documentation and Support Knowledge Base" on page 2

Release Notes

The Release Notes document, which is available on the SL Technical Support site at http://www.sl.com/support/, supplements the information in this user guide.

Documentation and Support Knowledge Base

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

Contacting SL

This section describes how to contact departments within SL.

Internet

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

Technical Support

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

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

CHAPTER 1 Quick Start

These instructions are for those customers who wish to evaluate the Monitor for purchase. These are the minimum steps required to gather monitoring data and get the Monitor up and running. Default settings are used and Apache Tomcat (which is delivered with the Monitor). These instructions are for both RTView Monitor for Solace AMI and On-premise version.

After you complete your evaluation, if you wish to setup and use all monitoring features in your organization, see "Configuration" (for both On-premise and AMI).

These instructions assume you have either the zip file for the On-premise option (either Windows or Linux) or a Solace AMI Instance. For details about Solace AMI Instances, see "Create Instance from RTView Monitor for Solace".

- **On-premise users**: See "Prerequisites for Solace On-premise Installations", then proceed to "On-premise Initial Steps".
- **AMI users**: See "Prerequisites for Solace AMI", proceed to "Create Instance from RTView Monitor for Solace" and then "Add Message Router Connections".

Prerequisites for Solace On-premise Installations

- Java JDK 1.7 or 1.8 64 bit
- Set the **JAVA_HOME** environment variable to point to your Java installation. For example:

  UNIX

  export JAVA_HOME=/opt/Java/jdk1.7.0

  Windows

  set JAVA_HOME=C:\Java\jdk1.7.0
- Linux On-premise Users:
  - These instructions require a Bourne-compatible shell.
  - JAVA_HOME is required to be in the PATH for Tomcat to start correctly.

You will need:

- Login credentials for each Solace message router you will monitor.

For complete RTView® system requirements, see **README_sysreq.txt**.

Proceed to "On-premise Initial Steps".

Prerequisites for Solace AMI

For complete RTView® system requirements, see **README_sysreq.txt**.

You will need:

- Login credentials for each Solace message router you will monitor.
- If you do not have an AMI instance, "Create Instance from RTView Monitor for Solace".

After you "Create Instance from RTView Monitor for Solace", proceed to "Add Message Router Connections".

On-premise Initial Steps

AMI users skip this step and proceed to "Add Message Router Connections".

1. Download **RTViewSolaceMonitor_<version>.zip** to your local server.
2. Extract the files:

   unzip -a RTViewSolaceMonitor_<version>.zip
3. Navigate to **RTViewSolaceMonitor/bin** directory and run **start_servers.sh** (**.bat** in Windows).

Proceed to "Add Message Router Connections," next.

Add Message Router Connections

These instructions describe how to connect your Solace Message Routers using the "RTView Configuration Application". These instructions apply to both On-premise and AMI versions.

1. Open a browser and type the following URL to open the RTView Configuration Application:
   - **http://IPAddress:8068/rtview/solmon_rtvadmin** if you are running the Monitor remotely.
   - **http://localhost:8068/rtview/solmon_rtvadmin** if you are running the Monitor locally.

   Use **rtvadmin**/**rtvadmin** as the username/password.

   The RTView Configuration Application **HOME** page opens.
2. Scroll down and select the **Solace Message Router Monitor** project.

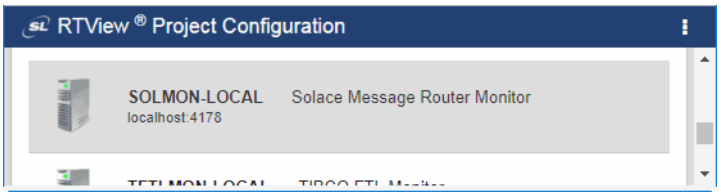

   The configuration page for the **Solace Message Router Monitor** project opens.

3. Select **Solace Message Router** in the navigation tree (left panel).

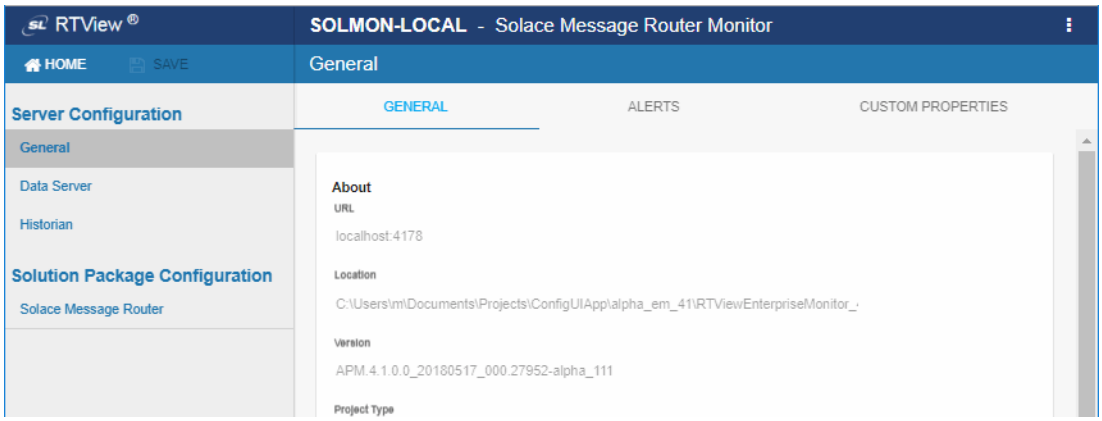

The **CONNECTIONS** tab opens.

4. Select **MESSAGE ROUTER** and click + to add a message router connection.

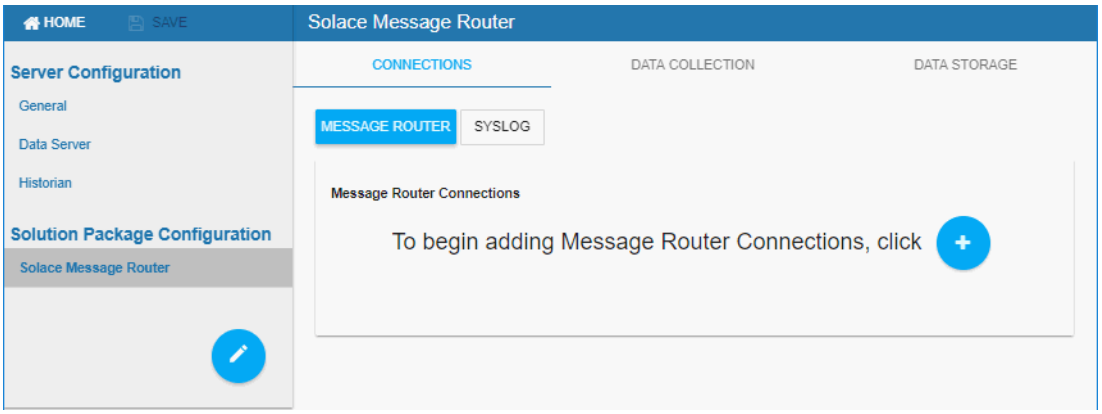

The **Add Connection** dialog opens.

5. In the **Add Connection** dialog: enter the **Connection Name**, **URL**, **Username** and **Password**, toggle on **Edition** if this is a VMR, enter the **SEMP Version** and the **VPN Name**.

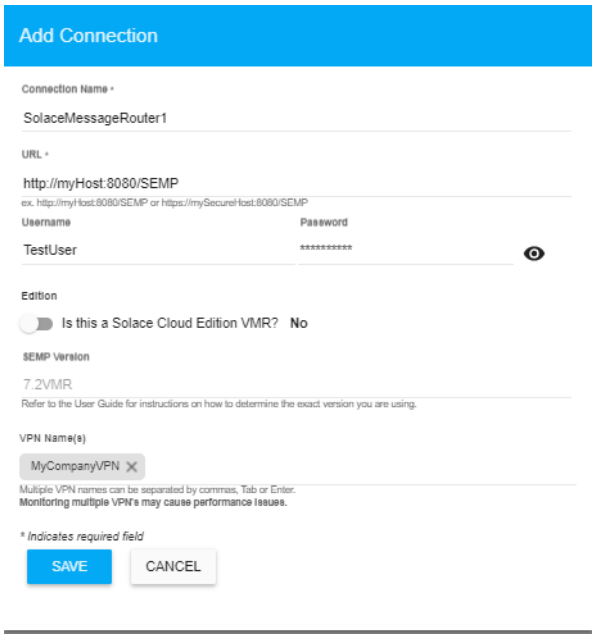

- **Notes:**
- if you want to collect additional metrics for a specific VPN, enter the **VPN Name**. Use this option carefully as this can increase the amount of data collected and impact monitor performance.
- if your message router is a Solace Cloud Edition VMR, toggle on the **Edition** option and enter the SEMP Version for previous versions to Solace 7.2. Otherwise, leave the default value.
- monitoring data is collected through SEMP (Solace Element Management Protocol), which is a RESTful API for managing Solace message routers. If your virtual message routers (VMRs) are a **version prior to 8.7 or Solace Appliance version prior to 8.3**) see "Obtain SEMP Version" for instructions on getting the SEMP version installed in your message routers.

6. Click SAVE to close the dialog and SAVE (in title bar) to save your settings.

The connections you create are listed in the **Connections** tab.

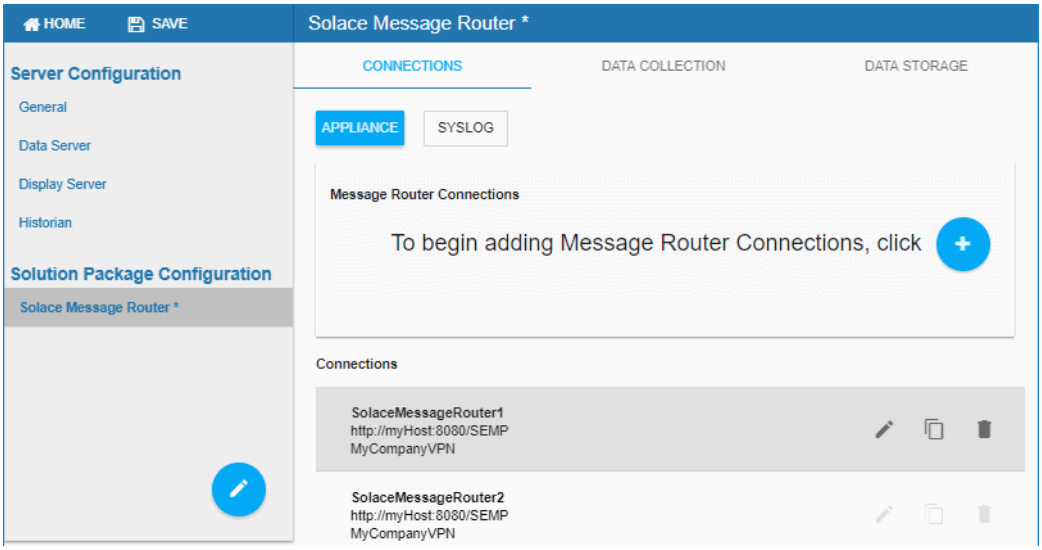

7. Repeat these steps to add multiple message routers and click SAVE when finished.

8. Click HOME to return to the **HOME** page, then click RESTART DATASERVER to apply your settings. The data server will be available again in 10-15 seconds. (Note that the **RESTART DATASERVER** option is visible only if you saved your connection entries.)

9. Open a browser and go to the RTView Solace Monitor:

http://IPAdress:8068/rtview/solmon if you are running the monitor remotely

http://localhost:8068/rtview/solmon if you are running the monitor locally

Use **solmon**/**solmonpw** for username/password (if login is enabled).

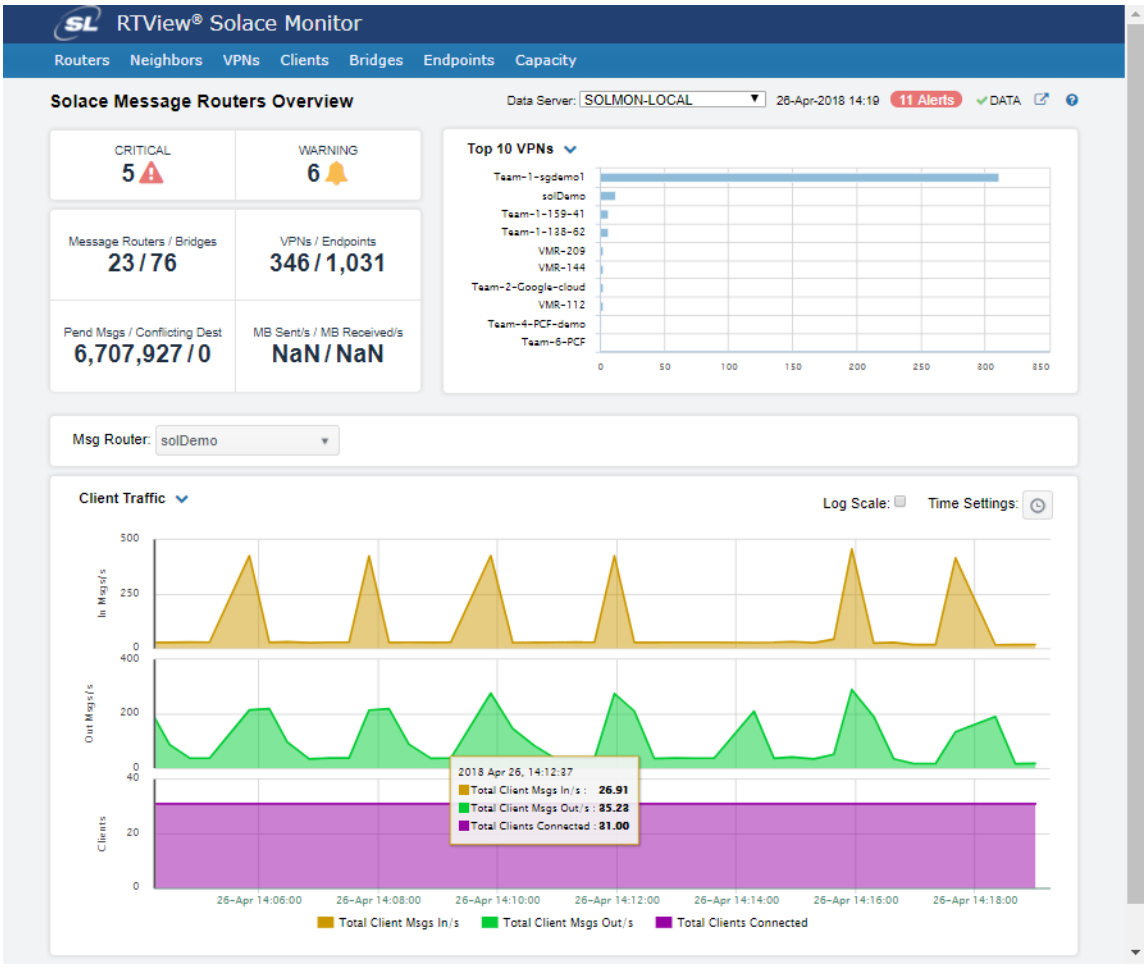

10. Verify that you see monitoring data. If you encounter issues, check the log files in the **RTViewSolaceMonitor/projects/rtview-server/log** directory for errors.

If you wish to setup and use all monitoring features in your organization, proceed to "Configuration" (for both On-premise and AMI).

You have completed the Quick Start!

CHAPTER 2 Introduction to the Monitor

This section contains the following:

- "Overview," next
- "System Requirements"
- "Installation"

Overview

The RTView Monitor for Solace is an easy to configure and use monitoring system that gives you extensive visibility into the health and performance of your Solace message routers and the infrastructure that relies on them.

The RTView Monitor for Solace enables Solace users to continually assess and analyze the health and performance of their infrastructure, gain early warning of issues with historical context, and effectively plan for capacity of their messaging system. It does so by aggregating and analyzing key performance metrics across all routers, bridges, endpoints and clients, and presents the results, in real time, through meaningful dashboards as data is collected.

Users also benefit from predefined dashboards and alerts that pin-point critical areas to monitor in most environments, and allow for customization of thresholds to let users fine-tune when alert events should be activated.

The RTView Monitor for Solace also contains alert management features so that the life cycle of an alert event can be managed to proper resolution. All of these features allow you to know exactly what is going on at any given point, analyze the historical trends of the key metrics, and respond to issues before they can degrade service levels in high-volume, high-transaction environments.

RTView Monitor for Solace is comprised of the following which you access with a browser:

- RTView Monitor for Solace, which monitors Solace performance metrics and used by teams to monitor the health of Solace components (message routers, bridges, clients, endpoints and VPNs). With the RTView Monitor for Solace AMI version, the health of "MySQL Database" and "Docker Engines" can be also monitored. For details, see "Using the Monitor".
- RTView Manager for Solace, which administrators use to set alert thresholds for RTView® Monitor for Solace®. For details, see "Using the Monitor".
- RTView Manager, which administrators use to monitor the health of RTView® Monitor for Solace®. That is, to monitor components of the monitoring system itself (RTView servers, JVMs, Tomcat servers, hosts, Docker, MySQL and alert settings for these components). For details, see "Using the Monitor".
- RTView Configuration Application, which administrators use to configure the majority of the monitoring system. For details, see "Configuration".

You can also install the monitor as a Solution Package (rather than a standalone product). See "Solution Package Version" for details.

RTView Monitor for Solace On-Premise Version

To evaluate the RTView Monitor for Solace, go to "Quick Start" to get up and running with RTView Monitor for Solace using default settings.

To install and use all features in RTView Monitor for Solace, see "Configuration".

RTView Monitor for Solace AMI Version

The RTView Monitor for Solace AMI version is an Amazon EC2 Amazon Machine Image (AMI) running Linux. It comes pre-configured with a 30-day license. The AMI instance includes an application stack including (among others) MySQL and Docker for convenience of quick deployment. Please refer to your instance's **/home/ec2-user/amibase/MANIFEST.txt** for the full version information.

To evaluate the RTView Monitor for Solace, go to "Quick Start" to get up and running with RTView Monitor for Solace using default settings.

To install and use all features in RTView Monitor for Solace, see "Configuration".

Solution Package Version

The RTView Monitor for Solace can also be installed as a Solution Package within the RTView Enterprise Monitor® product. RTView Enterprise Monitor is an end-to-end monitoring platform that allows application support teams to understand how infrastructure, middleware, and application performance data affect the availability and health of the entire system. Used as a Solution Package within RTView Enterprise Monitor, the Solace metrics are but one source of data, among many other sources (solution packages are available for TIBCO EMS, Amazon CloudWatch, TIBCO BusinessWorks, MicroSoft SQL and many others), that determine the entire health state of the application.

For more information about RTView Enterprise Monitor®, see the *RTView Enterprise Monitor® User's Guide*, available at http://www.sl.com/support/documentation/.

System Requirements

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

Installation

The Monitor can also be installed as a Solution Package within the RTView Enterprise Monitor® product.

Download the **RTViewSolaceMonitor_<version>.zip** file and unzip the **RTViewSolaceMonitor_<version>.zip** file into a directory of your choosing.

For more information about RTView Enterprise Monitor see the *RTView Enterprise Monitor® User's Guide*, available on the SL Product Documentation website.

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.

On UNIX/Linux systems, use the -a option to properly extract text files.

For more information about RTView Enterprise Monitor see the *RTView Enterprise Monitor® User's Guide*, available on the SL Product Documentation website.

CHAPTER 3 Configuration

This chapter describes how to change default settings and configure all features and components in the RTView Monitor for Solace (On-Premise and AMI). You will use the "RTView Configuration Application" to setup and modify the RTView Monitor for Solace *project*.

See "Quick Start" instructions if you want to evaluate the RTView Monitor for Solace.

Overview

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

- "Configure Data Collection": (**Required**) Define the Solace message routers to be monitored. This step must be performed before running any deployment of the Monitor. This section also describes how to configure history properties for storage and aggregation history properties for storage and aggregation of collected data, which is not required.
- "Change Port Assignments": (**Optional and for On-premise only**) Change the default port settings. This section does not apply to AMI.
- "Configure the Database": (**Optional and for On-premise only**) Configure a production database--for On-premise only. This section does not apply to AMI (AMI is pre-configured with MySQL).
- "Configure Alert Notifications": (**Optional**) Configure alerts to execute an automated action (for example, to send an email alert).
- "Troubleshooting": Investigate configuration issues.

Assumptions

This document assumes that you:

- verified "System Requirements".
- installed the Monitor per instructions in "Installation".
- AMI users have an AMI instance. If not, see "Create Instance from RTView Monitor for Solace" for details. (**Required for AMI only**)

Initialize a Command Prompt or Terminal Window

To start any RTView process (Data Server, Historian, and so forth), you must first initialize a command line window on the host. A Bash-compatible shell is required.

To initialize a command line window, execute the **rtvapm_init** script. For example:

Windows

Go to your Monitor installation directory and type:

rtvapm_init

UNIX

The script used to initialize a terminal window depends on whether you are in csh or rsh (e.g.Linux, Mac OS X). With a Bourne shell, open a terminal window, go to your Monitor installation directory and type:

. ./rtvapm_init.sh

Configure Data Collection

This section describes how to define connections for each of your Solace Message Routers and collect real-time data from them. We also describe settings for storing collected data in history.

If you don't have special requirements for running the monitor (such as running multiple data collectors in the same host), there is no need to cover the optional subsections. Consult Technical Support before modifying other configurations to avoid the circumstance of future upgrade issues. This section contains:

- "Open the RTView Monitor for Solace Project": Use the "RTView Configuration Application" to setup and modify the RTView Monitor for Solace project.
- "Define Solace Message Router and Syslog Connections": (**Required**) Define connection details for your message routers using the RTView Configuration Application for Solace.

  Note that the monitoring data is collected through SEMP (Solace Element Management Protocol), which is a RESTful API for managing Solace message routers. If your virtual message routers (VMRs) are a **version prior to 8.7 or Solace Appliance version prior to 8.3**, see "Obtain SEMP Version" for instructions on getting the SEMP version installed in your message routers.
- "Modify Default Settings for Storing Historical Data": (**Optional**) Describes the default settings, how to enable/disable storage of historical data and how to change the default settings.

Open the RTView Monitor for Solace Project

Open a browser, type the following URL and login (**rtvadmin**/**rtvadmin** is the username/password):

- **http://IPAddress:8068/rtview/solmon_rtvadmin** if you are running the Monitor remotely.
- **http://localhost:8068/rtview/solmon_rtvadmin** if you are running the Monitor locally.

The RTView Configuration Application **HOME** page opens. The **HOME** page is where you select the project you wish to configure and also to restart data servers to apply your settings.

Scroll down and select the **Solace Message Router Monitor** project.

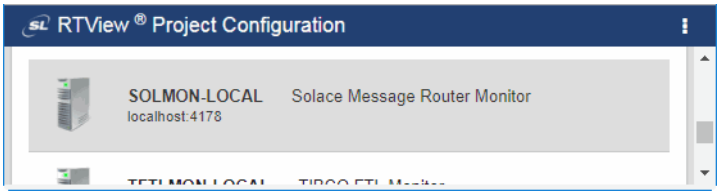

The configuration page for the **Solace Message Router Monitor** project opens.

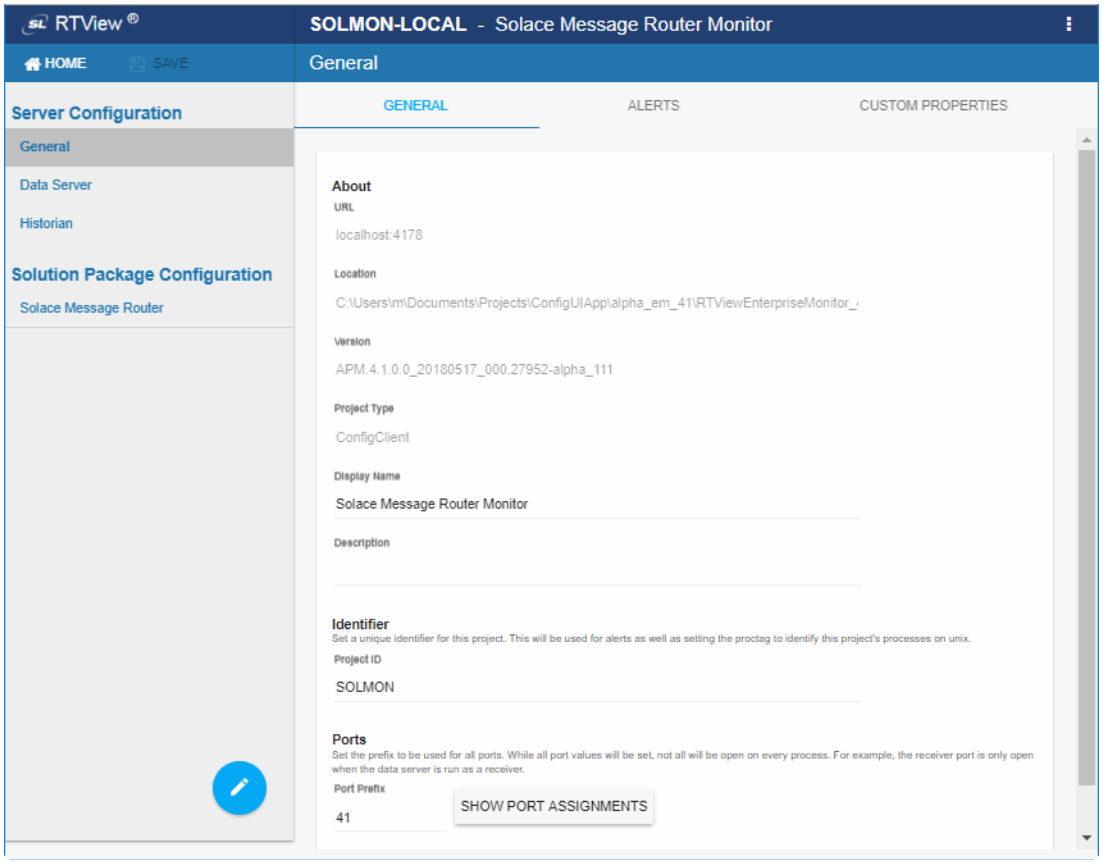

Tip: Use HOME (in title bar) to return to the **HOME** page.

Define Solace Message Router and Syslog Connections

You will define connections for the Solace message routers and Syslog you wish to monitor and verify you see collected data for these in the RTView Monitor for Solace.

1. "Open the RTView Monitor for Solace Project" and select **Solace Message Router** in the navigation tree (left panel). The following figure shows the top portion of the Solace Message Router configuration page.

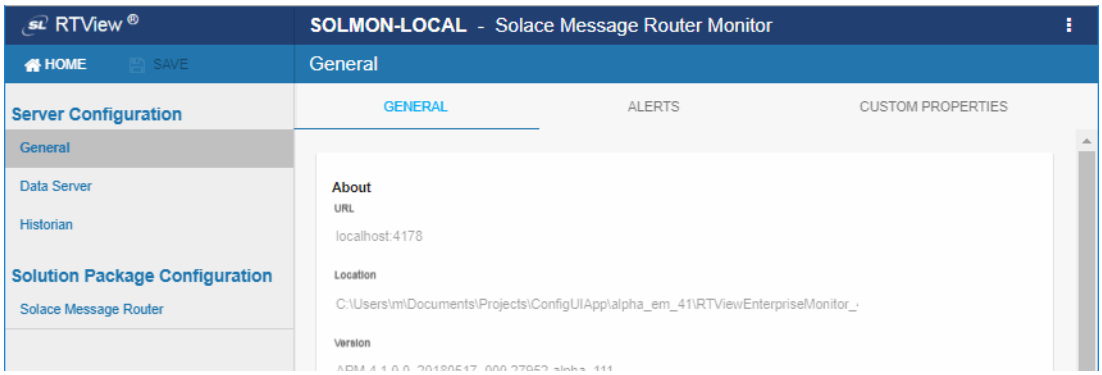

The **CONNECTIONS** tab opens.

2. In the **CONNECTIONS** tab, choose one of the following:

- Message Router: Choose **MESSAGE ROUTER** and click .
- Syslog: Choose **SYSLOG** and click .

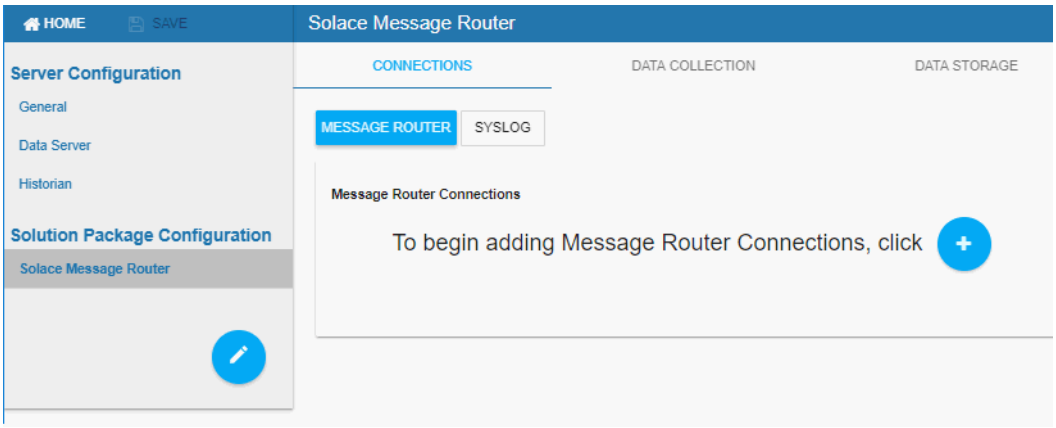

The **Add Connection** dialog opens.

3. In the **Add Connection** dialog, do one of the following. If you selected:

- **MESSAGE ROUTER**: Enter a **Connection Name**, **URL**, **Username** and **Password**, toggle on **Edition** if this is a VMR, enter the **SEMP Version** and the **VPN Name** and SAVE (in title bar) your entries.

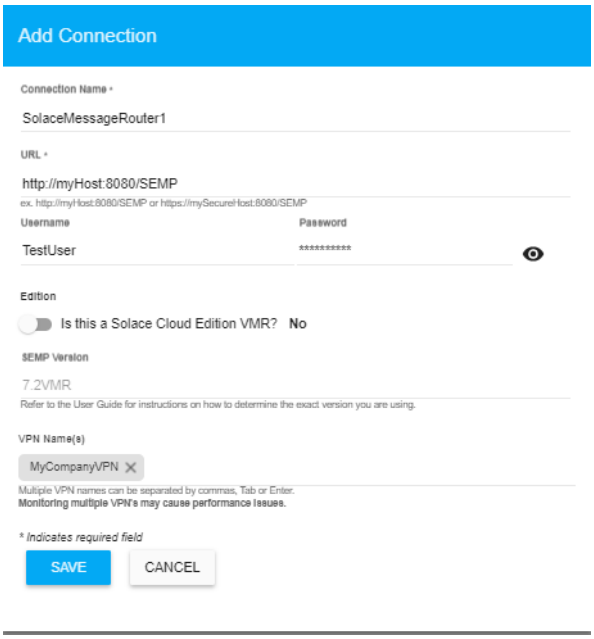

Notes:

- If you want to collect additional metrics for a specific VPN, enter the **VPN Name**. Use this option carefully as this can increase the amount of data collected and impact monitor performance.
- If your message router is a Solace Cloud Edition VMR, toggle on the **Edition** option and enter the SEMP Version for previous versions to Solace 7.2. Otherwise, leave the default value.

  - Monitoring data is collected through SEMP (Solace Element Management Protocol), which is a RESTful API for managing Solace message routers. If your virtual message routers (VMRs) are a **version prior to 8.7 or Solace Appliance version prior to 8.3**) see "Obtain SEMP Version" for instructions on getting the SEMP version installed in your message routers.

- **SYSLOG**: Enter a connection **Name**, choose either **TCP** or **UDP**, enter the **Host** name or the IP address of the local network interface, and enter the **Port** number from which incoming Syslog messages are read.

  By default, the TCP port is **601** and the UDP port is **514**.

Note: Only root can use ports **0** – **1024** on UNIX/Linux systems.

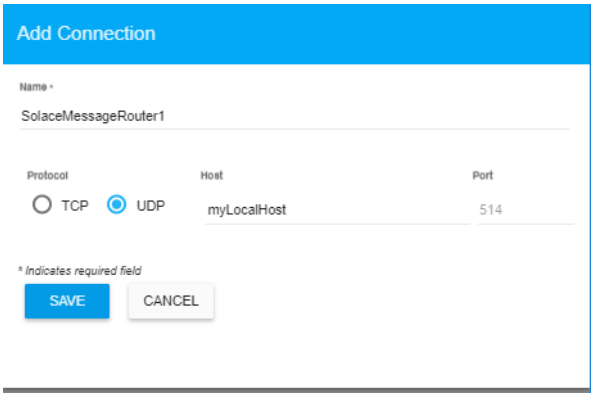

4. Click SAVE to close the dialog and SAVE (in title bar) to save your settings.

   The connections you create are listed in the **Connections** tab.

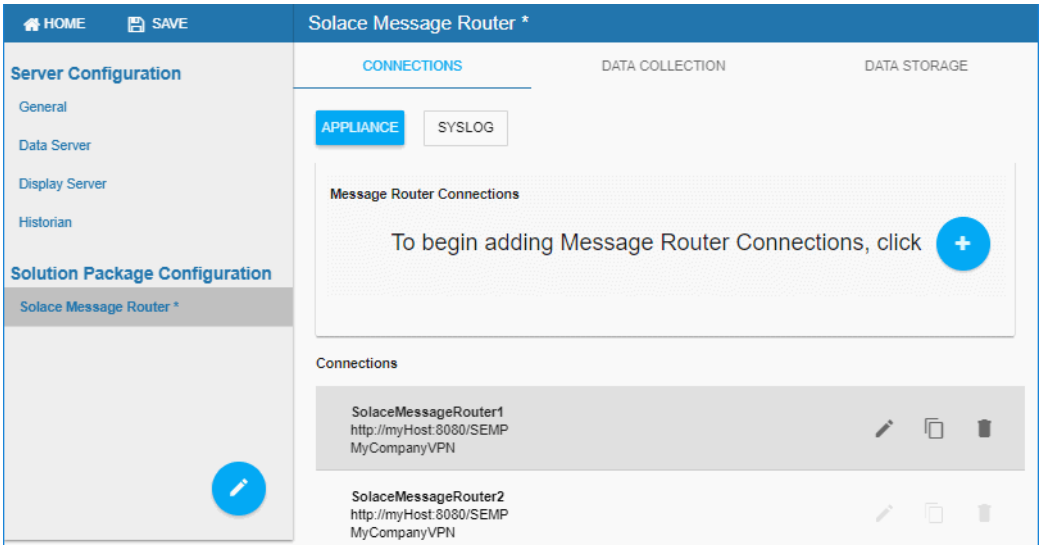

5. Repeat these steps using + to add as many connections as you need, then SAVE your settings.

6. In the **DATA COLLECTION** tab you can optionally modify the default polling rates for Solace caches.

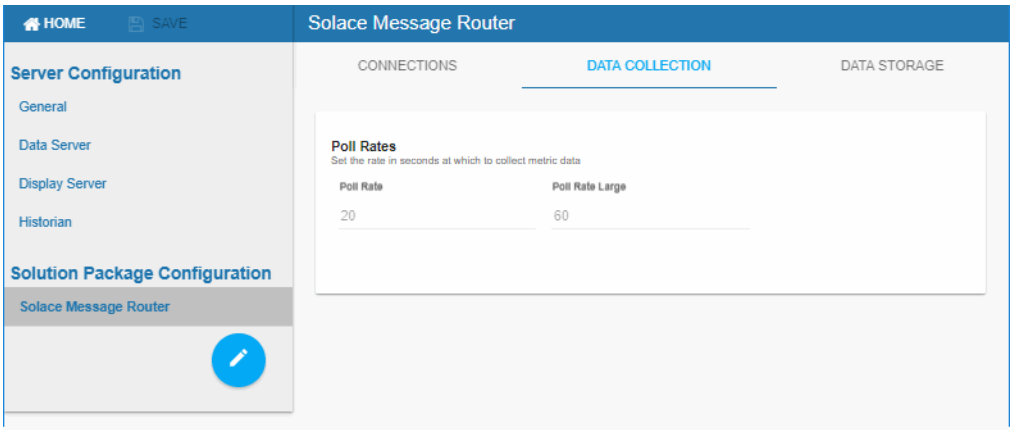

Poll Rates: Unit is seconds. Caches impacted are SolEndpointStats, SolEndpoints, SolClients, SolClientStats, SolBridges, SolAppliances, SolBridgeStats, SolApplianceInterfaces and SolApplianceMessageSpool.

Poll Rate Large: Caches impacted are SolCspfNeighbors, SolApplicances and SolEnvironmentSensors.

7. SAVE your settings, click HOME to return to the **HOME** page, then click RESTART DATASERVER to apply your settings. The data server will be available again in 10-15 seconds. (Note that the **RESTART DATASERVER** option is visible only if you saved your connection entries.)

8. Open a browser and go to the RTView Solace Monitor:

http://IPAddress:8068/rtview/solmon if you are running the monitor remotely

http://localhost:8068/rtview/solmon if you are running the monitor locally

Use **solmon**/**solmonpw** for username/password (if login is enabled).

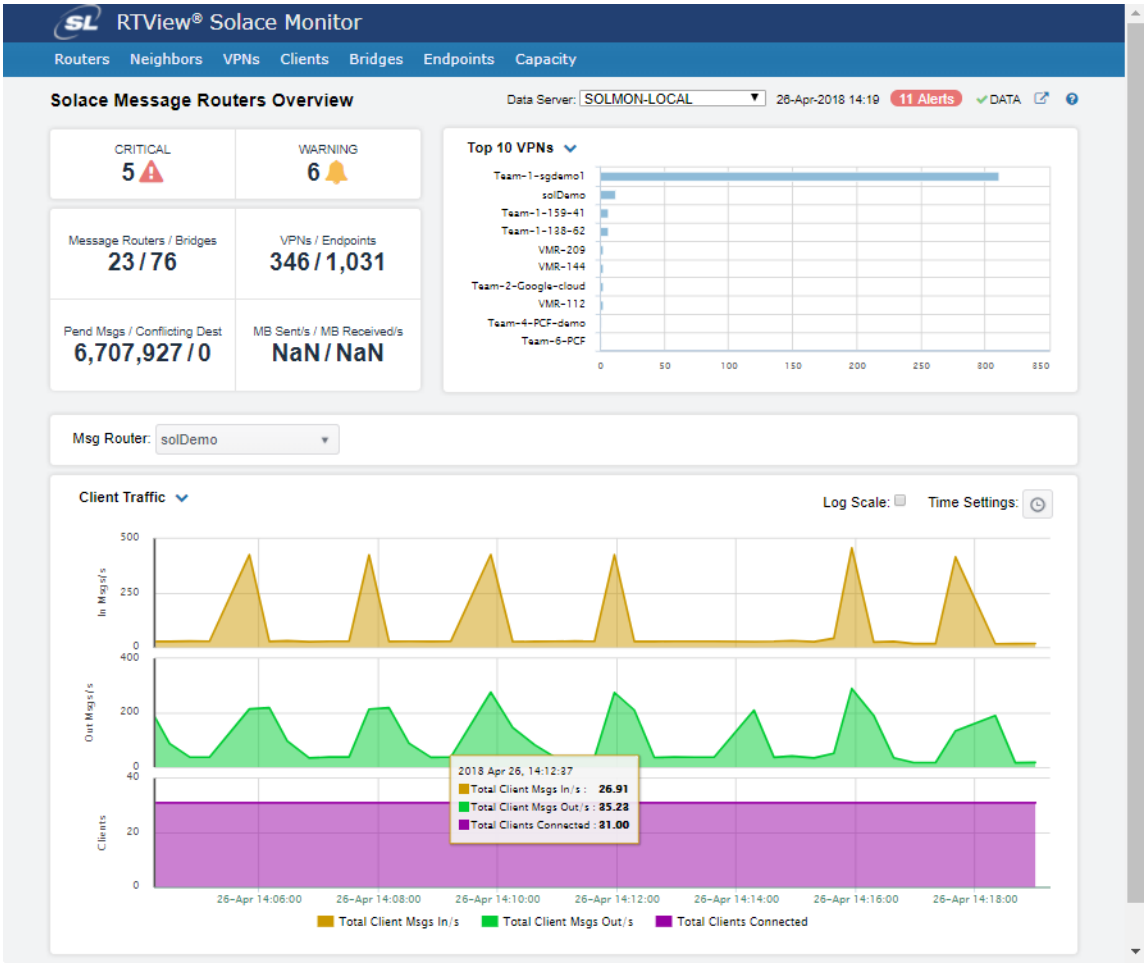

9. Verify that you see monitoring data. If you encounter issues, check the log files in the **RTViewSolaceMonitor/projects/rtview-server/log** directory for errors.

You have completed configuring data collection!

Optionally, On-premise users can proceed to:

- "Modify Default Settings for Storing Historical Data," next
- "Change Port Assignments"
- "Configure the Database"

Optionally, AMI and On-premise users can proceed to:

- "Configure Alert Notifications"
- "Troubleshooting"
- "Using the Monitor"

Modify Default Settings for Storing Historical Data

Use the RTView Configuration Application to change the default settings for storing historical data for Solace and the default cache settings that will modify the default behavior of the data being collected, aggregated and stored.

- "Enable/Disable Storage of Historical Data for Solace"
- "Define the Storage of In Memory Solace History"
- "Define Compaction Rules for Solace"
- "Define Expiration and Deletion Duration for Solace Metrics"
- "Enable/Disable Storage of Historical Data for Solace"
- "Define a Prefix for All History Table Names for Solace Metrics"

Enable/Disable Storage of Historical Data for Solace

Under **History Storage** you can select which tables you want the Historian to store in the database. To enable/disable the collection of historical data, perform the following steps:

1. "Open the RTView Monitor for Solace Project", select **Solace Message Router** in the navigation tree (left panel) and then select the **DATA STORAGE** tab (top right)..

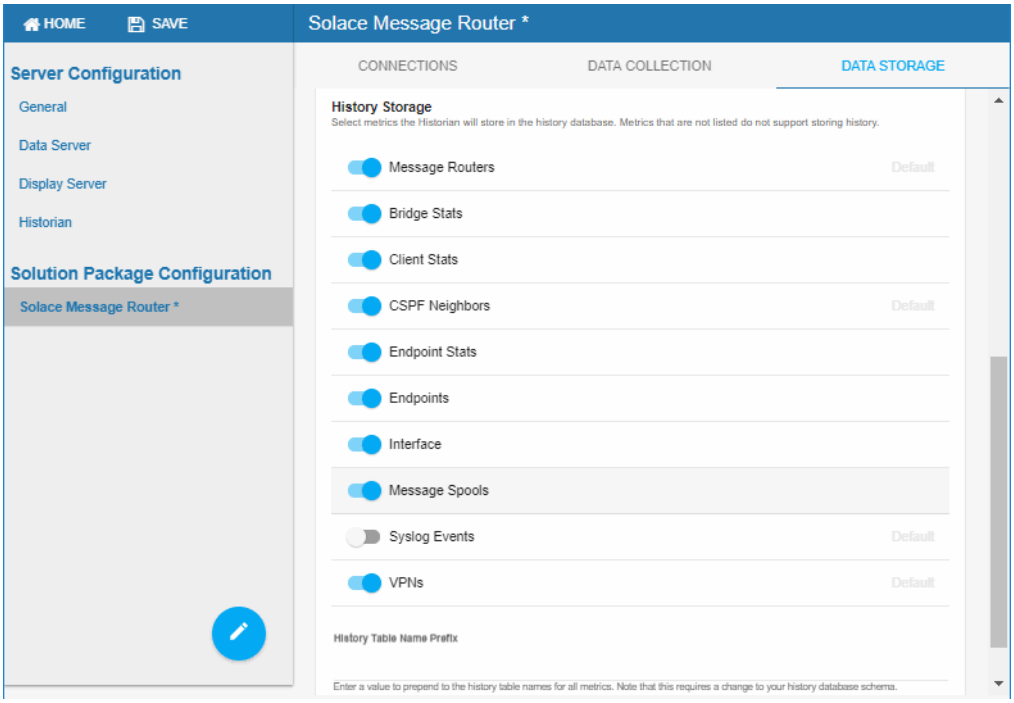

2. Scroll down to **History Storage** and toggle to enable/disable the storage of various database tables in the database. Blue (toggled right) enables storage, gray (toggled left) disables storage. The caches impacted by these settings are SolAppliances (**Message Routers**), SolBridgeStats (Bridge Stats), SolClientStats (Client Stats), SolCspfNeighbors (CSPF Neighbors), SolEndpointStats (Endpoint Stats), SolEndpoints (Endpoints),

SolApplianceInterfaces (**Interface**), SolApplianceMessageSpool (Message Spools), SolSyslogEvents (SyslogEvents) and SolVpns (VPNs).

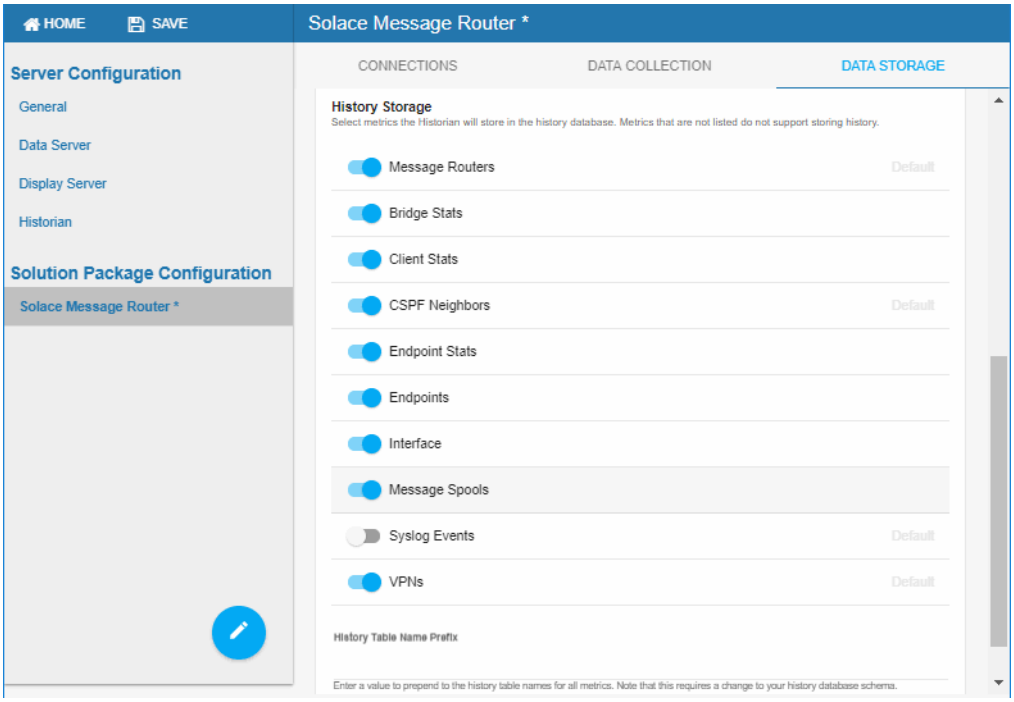

3. SAVE your settings, click HOME to return to the **HOME** page, then click RESTART DATASERVER to apply your settings.

Define the Storage of In Memory Solace 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 SolVpns, SolClientStats, SolAppliances, SolEndpoints, SolCspfNeighbors, SolBridgeStats, SolApplianceInterfaces, SolApplianceMessageSpool, SolEndpointStats and SolAppliancesQuality caches. The default setting for **History Rows** is **50,000**. To update the default settings:

To modify these settings do the following:

- "Open the RTView Monitor for Solace Project", select **Solace Message Router** in the navigation tree (left panel) and then select the **DATA STORAGE** tab (top right).
- In the **History Rows** field, enter the desired number of rows, SAVE your settings, click HOME to return to the **HOME** page, then click RESTART DATASERVER to apply your settings.

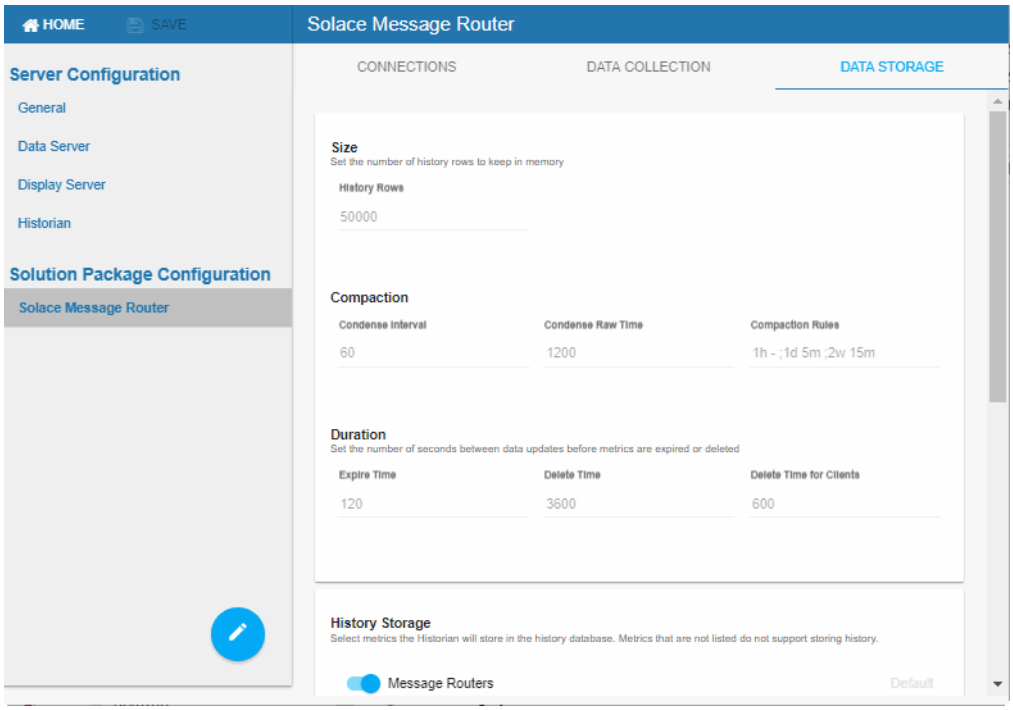

Define Compaction Rules for Solace

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. The default is 60 seconds. The following caches are impacted by this setting: SolVpns, SolClientStats, SolAppliances, SolEndpoints, SolCspfNeighbors, SolBridgeStats, SolApplianceInterfaces, SolApplianceMessageSpool and SolEndpointStats.
- **Condense Raw Time** -- The time span of raw data kept in the cache history table. The default is 1200 seconds. The following caches are impacted by this setting: SolVpns, SolClientStats, SolAppliances, SolEndpoints, SolCspfNeighbors, SolBridgeStats, SolApplianceInterfaces, SolApplianceMessageSpool and SolEndpointStats.
- **Compaction Rules** -- This field defines the rules used to condense your historical data in the database. 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). The following caches are impacted by this setting: SolVpns, SolClientStats, SolAppliances, SolEndpoints, SolCspfNeighbors, SolBridgeStats, SolApplianceInterfaces, SolApplianceMessageSpool and SolEndpointStats.

To modify these settings do the following:

- "Open the RTView Monitor for Solace Project", select **Solace Message Router** in the navigation tree (left panel) and then select the **DATA STORAGE** tab (top right).
- Enter values in the **Condense Interval**, **Condense Raw Time** and **Compaction Rules** fields, SAVE your settings, click HOME to return to the **HOME** page, then click RESTART DATASERVER to apply your settings.

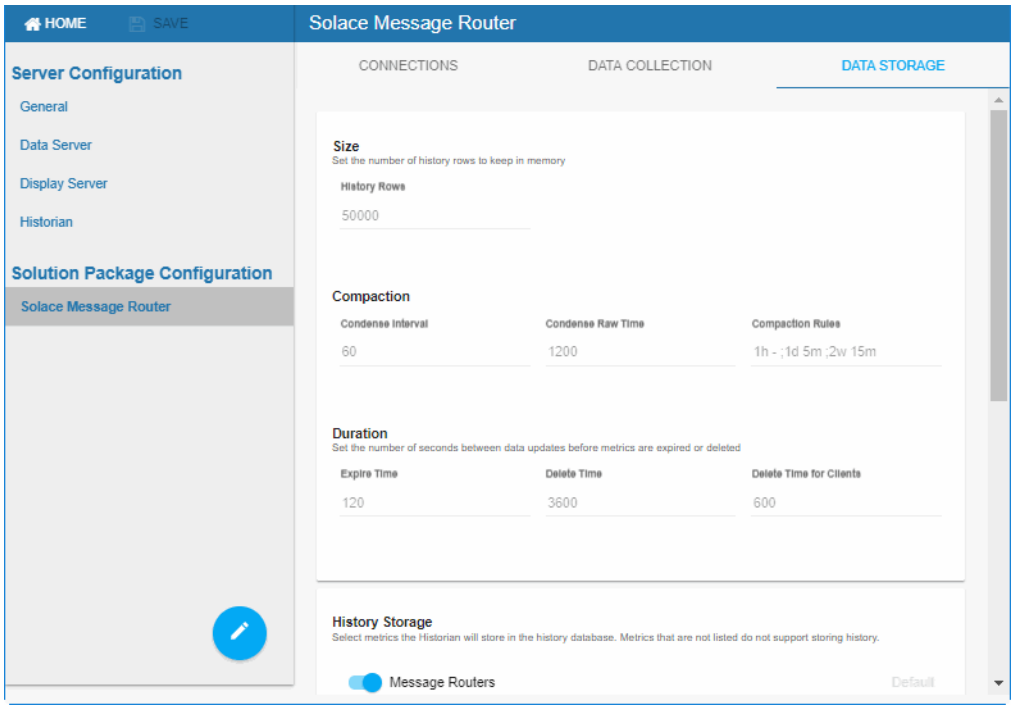

Define Expiration and Deletion Duration for Solace Metrics

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.

The **Expire Time** field, which sets the expire time for the SolVpns, SolBridges, SolClients, SolClientStats, SolAppliances, SolEndpoints, SolCspfNeighbors, SolBridgeStats, SolApplianceInterfaces, SolApplianceMessageSpool, SolEndpointStats, SolEnvironmentSensors and SolAppliancesQuality caches, defaults to 120 seconds.

The **Delete Time**, which sets the delete time for the SolVpns, SolBridges, SolEndpoints, SolBridgeStats, SolEndpointStat and SolEnvironmentSensors caches, defaults to 3600 seconds. To modify these defaults:

The **Delete Time for Clients**, which sets the delete time for the SolClients and SolClientStats caches, defaults to 600 seconds.

To modify these settings do the following:

- "Open the RTView Monitor for Solace Project", select **Solace Message Router** in the navigation tree (left panel) and then select the **DATA STORAGE** tab (top right).
- In the **Delete Time for Clients** field, enter the desired value, SAVE your settings, click HOME to return to the **HOME** page, then click RESTART DATASERVER to apply your settings.

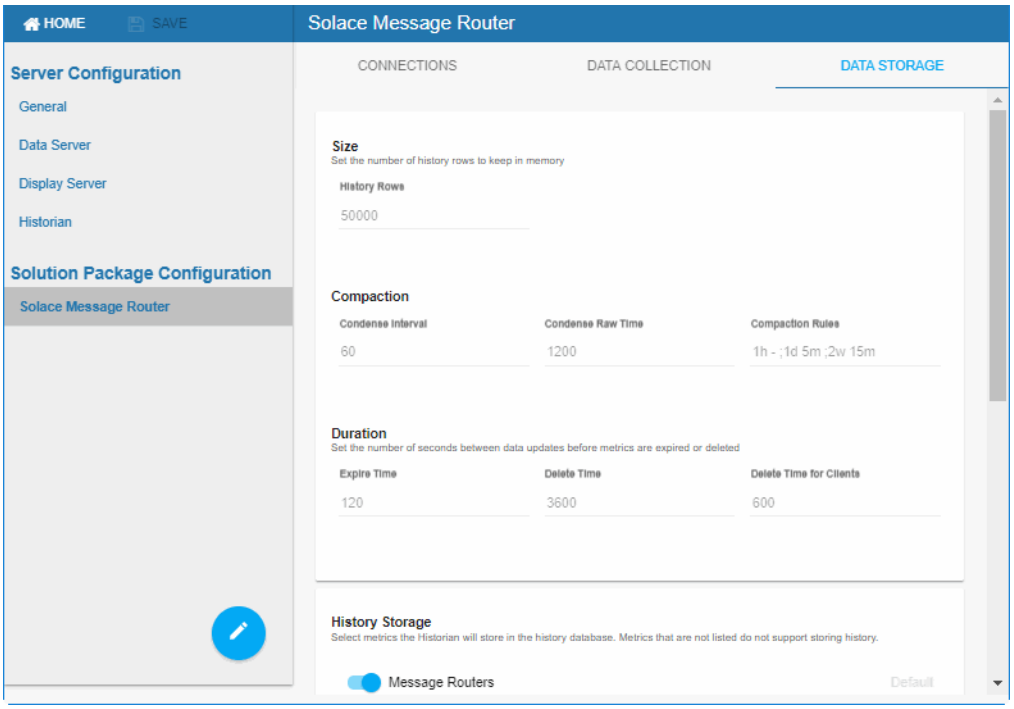

Define a Prefix for All History Table Names for Solace Metrics

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 **RTVAPM_HOME/solmon/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

To add the prefix do the following:

- "Open the RTView Monitor for Solace Project", select **Solace Message Router** in the navigation tree (left panel), select the **DATA STORAGE** tab (top right) and scroll down to the bottom of the page.
- In the **History Table Name Prefix** field, enter the desired prefix name, SAVE your settings, click HOME to return to the **HOME** page, then click RESTART DATASERVER to apply your settings.

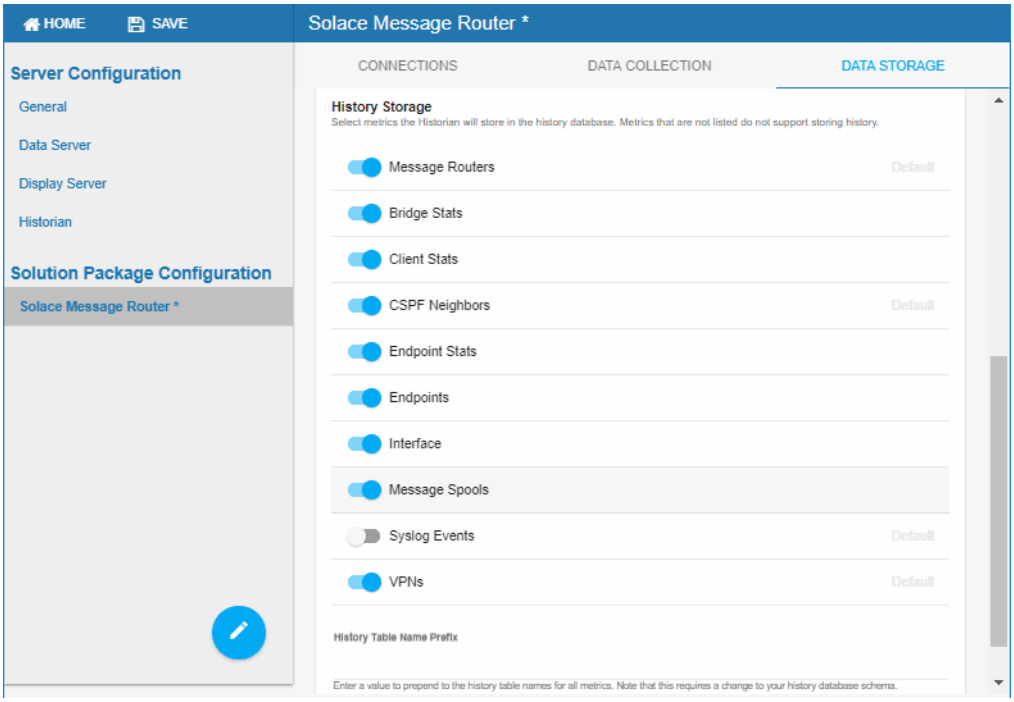

Change Port Assignments

This configuration is optional for the On-premise version. Ports should not be changed in the AMI version.

There are deployment architectures that might require the change of default ports for selected processes, either because the process will be executed multiple times in the same host or because the selected port number is already in use by another application. In these circumstances, you should reassign ports for Solace using the RTView Configuration Application.

Java Process	Description	Default Port(s)
RTView Data Server	Gathers performance metrics.	Default Port= **4178** Default JMX Port = **4168**

Receiver RTView Data Server	Receiver Data Server in a fault tolerant pair.	Default Port= **4172** Default JMX Port= **4168**
Sender RTView Data Server	Sender Data Server in a fault tolerant pair.	Default Port= **4176** Default JMX Port= **4166**
RTView Historian	Retrieves data from the RTView Data Server and archives metric history to a database.	Default JMX Port= **4167**
RTView Display Server	Collects the data and generates the displays that the Application Server uses to produce the web pages.	Default Port= **4179** Default JMX Port = **4169**

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

1. "Open the RTView Monitor for Solace Project", select **General** in the navigation tree (left panel). The **GENERAL** tab opens by default.

2. Under **Ports** (lower portion of main panel), specify the port prefix that you want to use in the **Port Prefix** field. Click **Show Port Assignments** to see the port numbers that are created using the **Port Prefix** you specify.

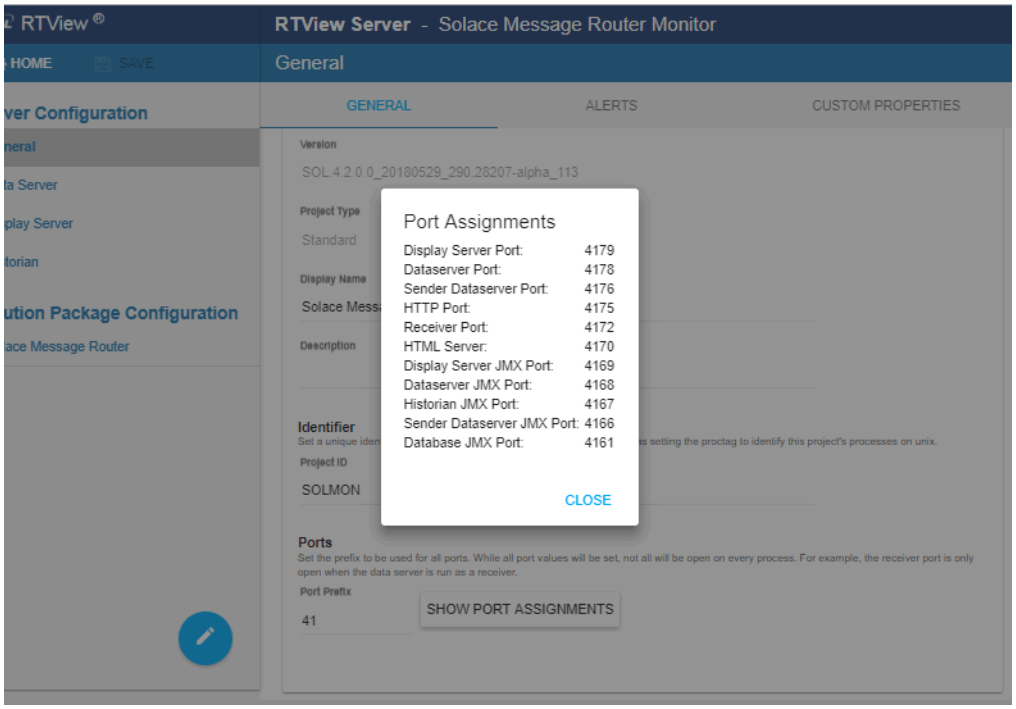

3. Click SAVE (in the title bar).

4. Restart all servers by running **RTViewSolaceMonitor/bin/stop_servers** (.bat or .sh) and then running **RTViewMonitor/bin/start_servers** (.bat or .sh).

5. Edit the **update_wars** (**.bat** or **.sh**) file and change the port prefix for all ports to the prefix you just specified.

6. Rebuild the war files and install them to the application server by executing the following script, located in the **RTViewSolaceMonitor/bin** directory:

   Windows:

   make_all.bat

   UNIX:

   ./make_all.sh

Configure the Database

This section is optional for RTView Monitor for Solace On-Premise. This section does not apply to the AMI version. This section describes how to setup an alternate (and supported) 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 *RTView Core® User's Guide*.

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, Sybase, 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 in the "Define the ALERTDEFS DB" section.

3. "Open the RTView Monitor for Solace Project", select **General** in the navigation tree (left panel), then the **ALERTS** tab and click the **Edit** icon (the light gray pencil icon to the right of) **Alert Threshold Database Connection**.

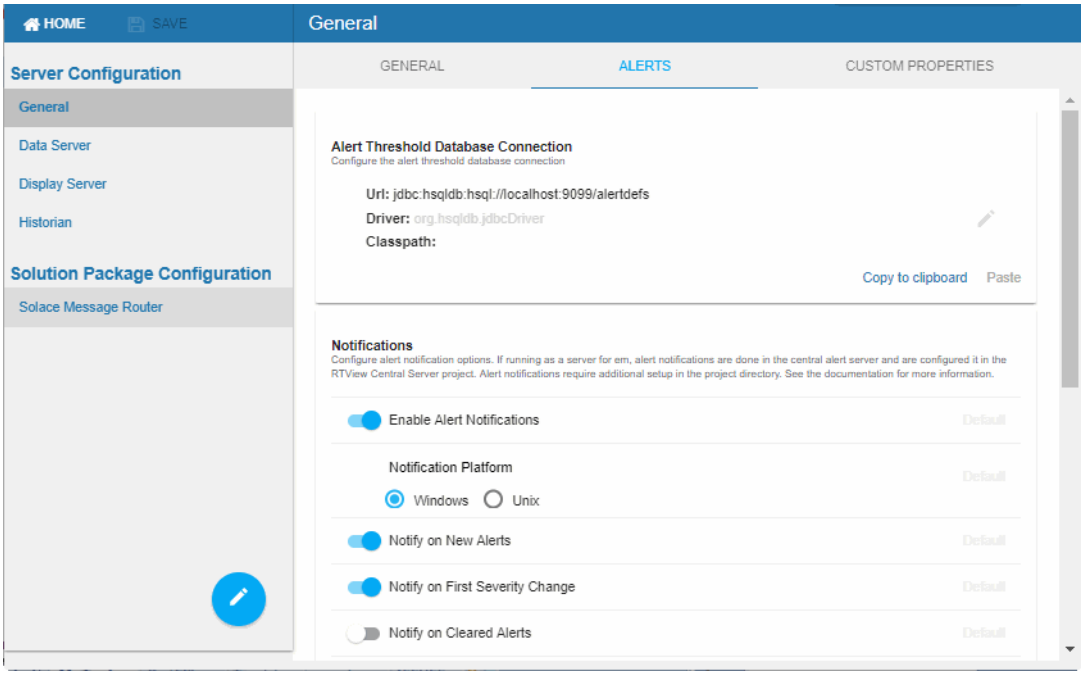

The **Edit Connection** dialog opens.

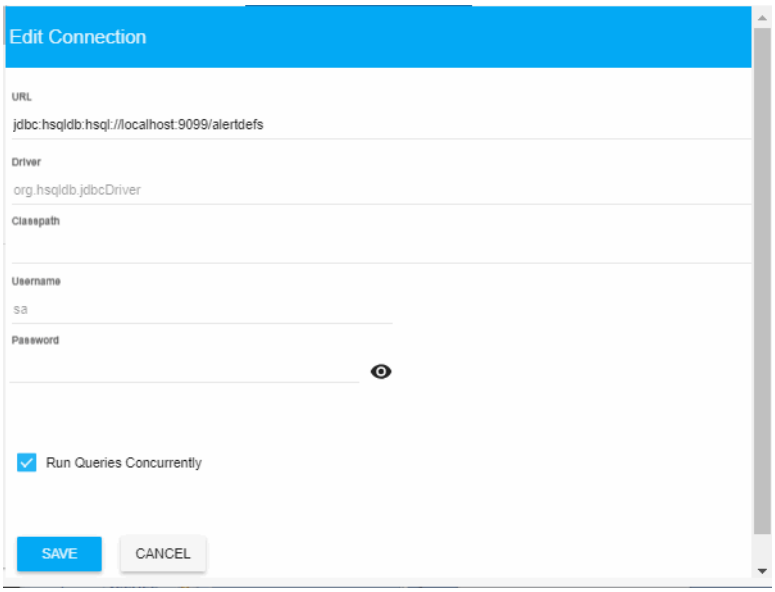

4. Enter the information from Step 2 into the **Edit Connection** dialog and click **Save**.

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.

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

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

6. In the RTView Configuration Application, select **Historian** in the navigation tree (left panel) and click the **Edit** icon (light gray pencil to the right of) **Historian Database Connection** (in the main panel).

The **Edit Connection** dialog opens.

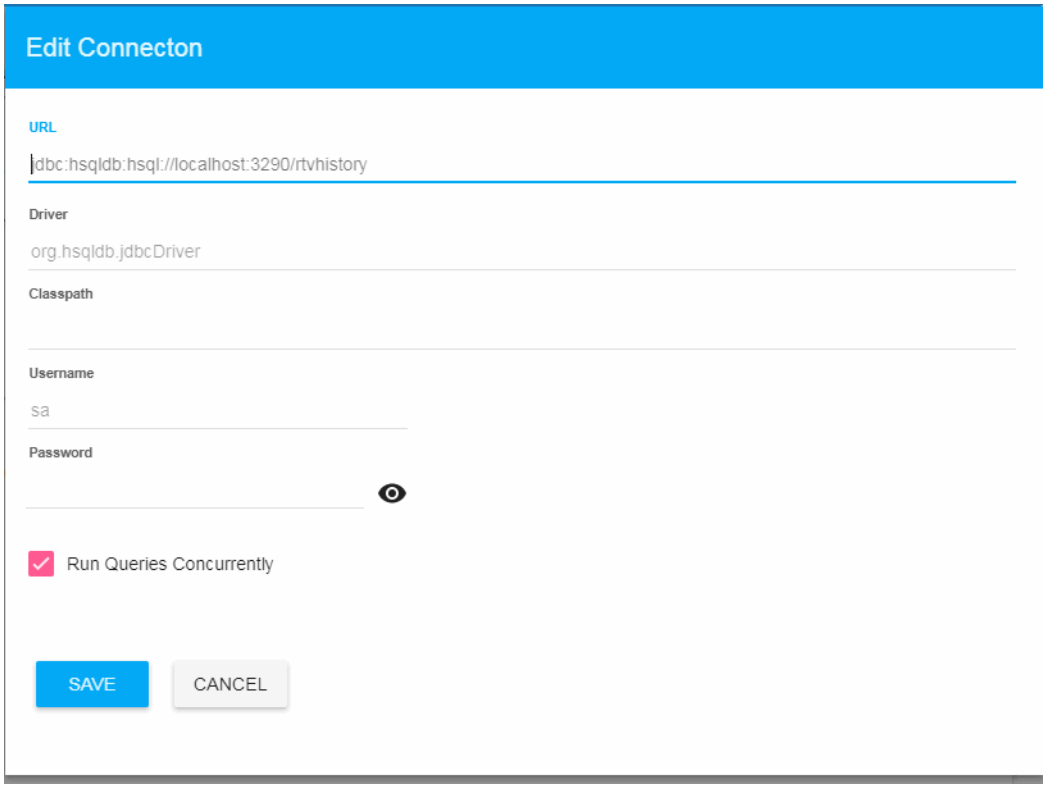

7. Enter the information from Step 5 into the **Edit Connection** dialog and click **Save**.

   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.

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

8. Click in the RTView Configuration Application title bar.

9. Restart all servers by running **RTViewSolaceMonitor/bin/stop_servers** (.bat or .sh) and then running **RTViewMonitor/bin/start_servers** (.bat or .sh).

10. 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** and **solmon** directories:

- **Alerts**

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

- **History**

  rtvapm/solmon/dbconfig/create_solmon_history_tables_<db>.sql

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

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. Proceed to Configure Alert Notification.

Configure Alert Notifications

The Monitor provides alerts concerning conditions in your system through RTView alerts. This section describes how to configure the alerts to execute an automated action.

You can configure alerts to notify on the following events:

- When a new alert is created
- The first time the **Severity** field on an alert changes
- When an alert is cleared
- Periodically renotify for unacknowledged alerts

By default, a **.bat** script is executed for new alerts, as well as when the first severity change occurs 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,
ALERTINDEX=alertIndex1~alertIndex2, ALERTID=1075, ALERTSEVERITY=2, ALERTTEXT=High Alert
Limit exceeded current value: 100.0 limit: 80.0 #####
```

To configure alert notifications, you need to:

1. Configure when to execute alert notifications and what action to perform using the RTView Configuration Application. See "Configuring Alert Notifications using the RTView Configuration Application" for more information.
2. Configure either of the two options for alert notification actions (Batch File/Shell Script or Java Command Handler). See "Configuring Monitor Alert Notification Actions" for more information.

Configuring Alert Notifications using the RTView Configuration Application

1. "Open the RTView Monitor for Solace Project", select **General** in the navigation tree (left panel) and then select the **ALERTS** tab (in the main panel)**.**

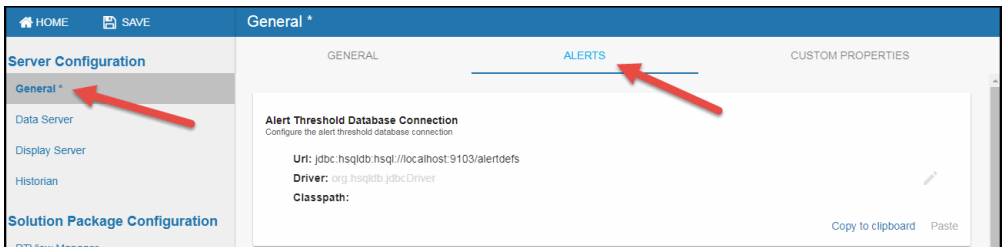

2. Under **Notifications**, toggle on **Enable Alert Notifications** (toggle is blue in ON position). If you don't see the **Enable Alert Notifications** option, toggle on **Configure notifications for this server in addition to central notifications** to expand the options.

3. If you are going to execute a script for your alert notifications, select the proper option (**Windows**/**Unix**) for **Notification Platform** to specify in which platform the project is running.

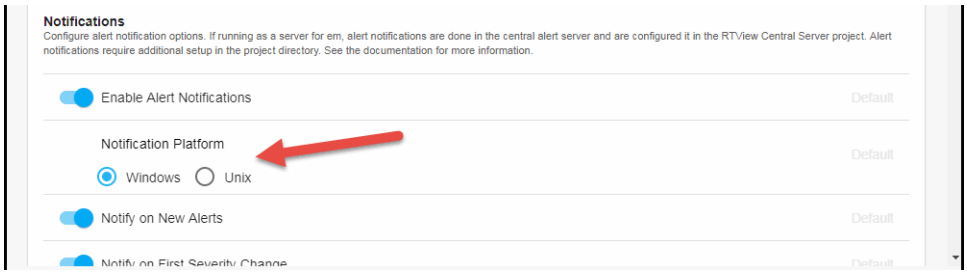

4. Select the events on which you would like to be notified (blue is enabled/gray is disabled):

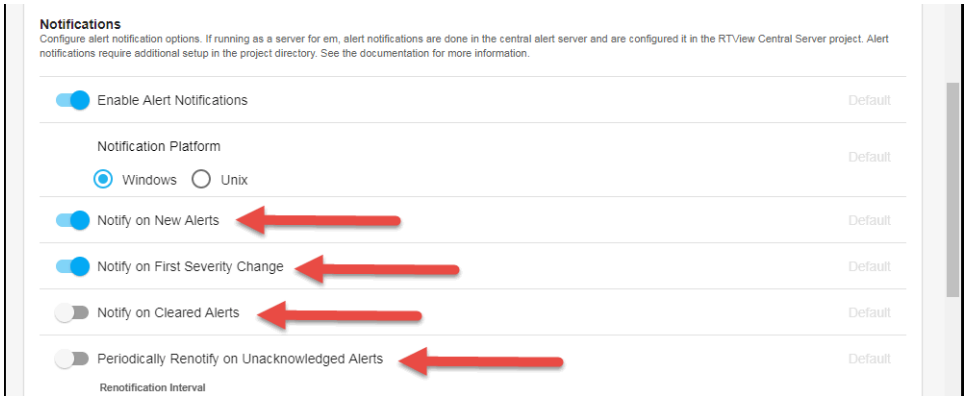

Notify on New Alerts: your action will be executed every time a new alert is created.

Notify on First Severity Change: your action will be executed the first time the severity changes for each alert.

Notify on Cleared Alerts: your action will be executed every time an alert is cleared.

Periodically Renotify on Unacknowledged Alerts: your action will be executed on the Renotification Interval (in seconds) for each unacknowledged alert.

5. If you will be executing a script for your alert notifications, skip to **Step 7**. If you will be executing the Java command: select **General** in the navigation tree (left panel), select the **CUSTOM PROPERTIES** tab and click + . The **Add Property** dialog opens.

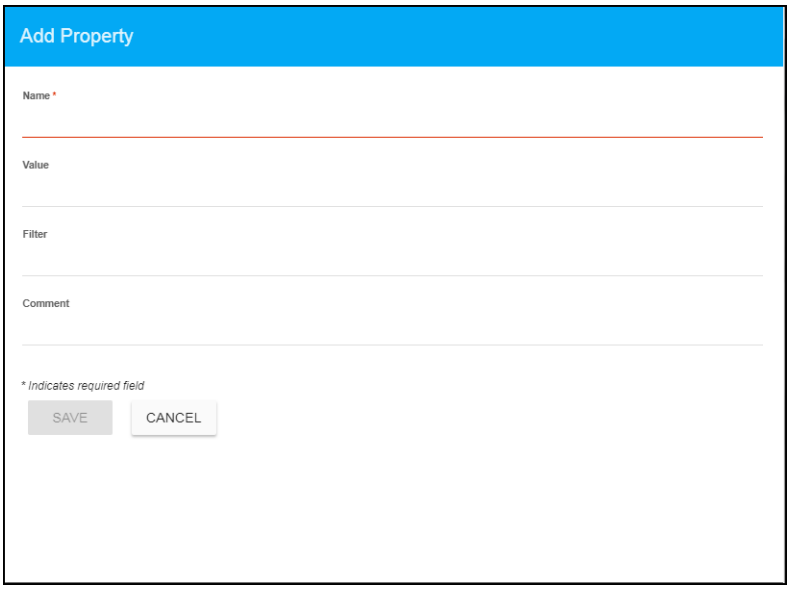

6. Create the following custom properties, one at a time, and click **SAVE** after creating each:

Name: sl.rtview.cp

Value: ./custom/lib/rtvapm_custom.jar

Filter: dataserver

Name: sl.rtvapm.customcommandhandler

Value: com.sl.rtvapm.custom.RtvApmCommandHandler

Filter: dataserver

If you selected the **Notify on New Alerts** option in Step 4, add:

Name: sl.rtview.alert.notifiercommandnew

Value: system.cust
'my_alert_notification.$domainName.$alertNotifyType.$alertNotifyCol' $alertNotifyTable

Filter: dataserver

If you selected the **Notify on First Severity Change** option in Step 4, add:

Name: sl.rtview.alert.notifiercommandfirstsevchange

Value: system.cust
'my_alert_notification.$domainName.$alertNotifyType.$alertNotifyCol' $alertNotifyTable

Filter: dataserver

If you selected the **Notify on Cleared Alerts** option in Step 4, add:

Name: sl.rtview.alert.notifiercommandcleared

Value: system.cust ‘my_alert_notification.$domainName.$alertNotifyType.$alertNotifyCol’ $alertNotifyTable

Filter: dataserver

If you selected the **Periodically Renotify on Unacknowledged Alerts** option in Step 4, add:

Name: sl.rtview.alert.notifiercommandrenot

Value: system.cust ‘my_alert_notification.$domainName.$alertNotifyType.$alertNotifyCol’ $alertNotifyTable

Filter: dataserver

After all three are created and saved, the newly created properties display in the **Custom Properties** tab.

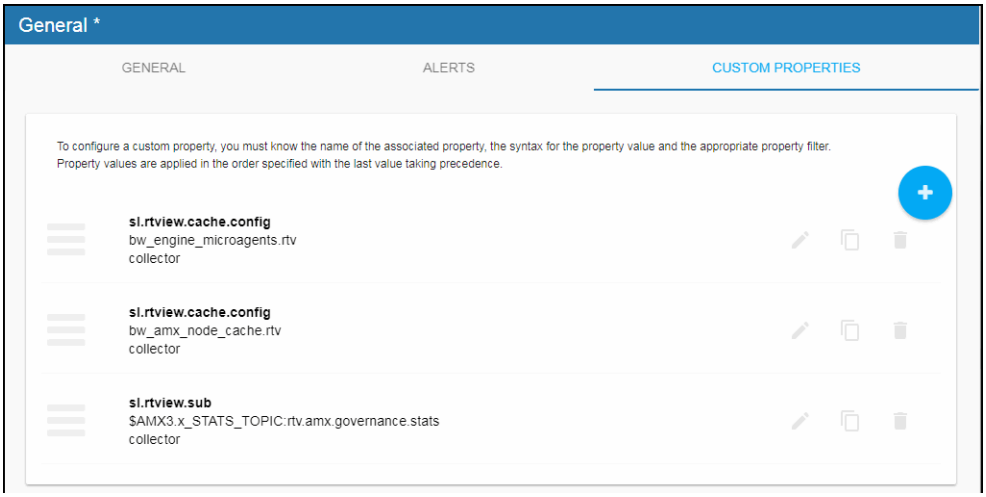

7. Click SAVE and restart the data server (after you have configured your monitor alert notifications below) to apply your changes.

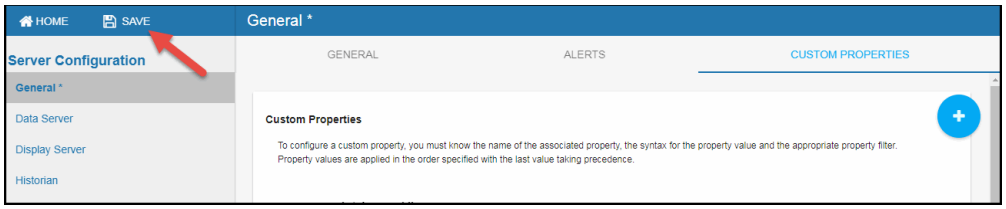

8. Configure Monitor alert notification actions via batch file/shell script or via the Java Command Handler. See the specific steps for each in the “Configuring Monitor Alert Notification Actions” section below.

Configuring Monitor Alert Notification Actions

There are two options for configuring Monitor alert notification actions:

- "Using a Batch File or Shell Script" on page 35:This technique requires switching to an OS-specific set of alert definitions that execute the appropriate file type. Windows and UNIX alert definition files are provided with the Monitor. A sample batch file and a sample shell script are also provided which are customized as needed.
- "Using the Java Command Handler" on page 37:The Java source for the Monitor Java command handler is provided to facilitate customization.

Using a Batch File or Shell Script

A sample batch file, **my_alert_actions.bat**, and a sample shell script, **my_alert_actions.sh**, which are located in the **common/bin** directory, are provided as templates that you can modify as needed. Use the appropriate file for the platform that hosts Monitor processes. By default, both scripts send alert information to standard output. To uncomment the line in the script so that alert data prints to standard output in:

- "Windows Batch File," next
- "UNIX/Linux Shell Script" on page 35

Windows Batch File

1. Copy the **my_alert_actions.bat** file, located in the **common/bin** directory, into your project directory.

2. Open the **my_alert_actions.bat** file, located in your project directory, and uncomment the echo line (near the end of the file) to print alert information to standard output. Or, you can modify the script to execute an automated action (such as sending an email alert). This script will be executed for new alerts and on first severity change.

3. If you selected **Notify on Cleared Alerts** in the RTView Configuration Application, copy **my_alert_actions.bat** from step 2 to **my_alert_actions.cleared.bat**. Optionally modify the script to execute a different action for cleared alerts. This script will execute when an alert is cleared.

4. If you selected **Periodically Renotify on Unacknowledged Alerts** in the RTView Configuration Application, copy **my_alert_actions.bat** from step 2 to **my_alert_actions.renotify.bat**. Optionally modify the script to execute a different action for renotifications. This script will execute periodically for unacknowledged alerts.

5. Restart the Data Server.

UNIX/Linux Shell Script

1. Copy the **my_alert_actions.sh** file, located in the **common/bin** directory, into your project directory.

2. Open the **my_alert_actions.sh** file, located in your project directory, and uncomment the echo line (near the end of the file) to print alert information to standard output. Or, you can modify the script to execute an automated action (such as sending an email alert). This script will be executed for new alerts and on first severity change.

3. If you selected **Notify on Cleared Alerts** in the RTView Configuration Application, copy **my_alert_actions.sh** from step 2 to **my_alert_actions.cleared.sh**. Optionally modify the script to execute a different action for cleared alerts. This script will execute when an alert is cleared.

4. If you selected **Periodically Renotify on Unacknowledged Alerts** in the RTView Configuration Application, copy **my_alert_actions.sh** from step 2 to **my_alert_actions.renotify.sh**. Optionally modify the script to execute a different action for renotifications. This script will execute periodically for unacknowledged alerts.

5. Restart the Data Server.

Batch File or Shell Script Substitutions

The default **my_alert_actions** scripts use the substitutions described in the table below.

Argument	Description	Values
$alertId	This substitution specifies the unique ID for the alert. For example: **alertId = 1004**	Text or Numeric
$alertIndex	This substitution specifies which source triggered the alert. With tabular objects, the first column of data is typically the **Index** column. The value in the **Index** column is a name that uniquely identifies each table row. The **alertIndex** uses the **Index** column name. For example, if the **CapactityLimitAllCaches** alert is configured to monitor all of your caches, and to trigger when any of the caches exceed the specified capacity threshold, the **alertIndex** indicates specifically which cache triggered the alert. With scalar objects, which do not have a table and therefore do not have a column (the **useTabularDataFlag** property is **False**), the **alertIndex** is blank. For example: **alertIndex = MyCache01**	Text or Numeric
$alertName	This substitution specifies the name of the alert. For example: **alertName = CapacityLimitAllCaches**	Values vary.
$alertSeverity	This substitution specifies the severity level of the alert. **0**: The alert limit has not been exceeded therefore the alert is not activated. **1**: The alert warning limit has been exceeded. **2**: The alert alarm limit has been exceeded. For example: **alertSeverity = 1**	Numeric
$alertText	This substitution specifies the text that is displayed when the alert executes. For example: **alertText = High Warning Limit exceeded, current value: 0.9452 limit: 0.8**	Text

Using the Java Command Handler

1. Verify that the **rtvapm_custom.jar** file is built per Step 4 in the "Customizing the Java Command Handler" instructions.

2. Restart the Data Server.

Customizing the Java Command Handler

The source for the Monitor Java handler is provided in the **RtvApmCommandHandler.java** file, located in the **\projects\custom\src\com\sl\rtvapm\custom** directory of your Monitor installation 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_classes.bat** or **make_classes.sh**) in the **\projects\custom\src** directory.

5. Restart the Data Server.

Java Command Handler Substitutions

When you customize the Java Command Handler, the entire alert table row is passed into the Java Command Handler for each alert that notifies so that all information regarding those alerts is available. The following substitutions are used:

Argument Description

- **$alertNotifyType -** This substitution specifies to show the value of the notification type so you can use the same command for all notifications. Values are **NEW_ALERT**, **CLEARED_ALERT**, **FIRST_SEV_CHANGE** or **COLUMN_CHANGED**.
- **$alertNotifyCol -** This substitution only applies when the **notifyType** is **COLUMN_CHANGED**. Specifies to use a semi-colon delimited list of column names that changed from the **alertNotifierColumns**.
- **$alertNotifyTable -** This substitution specifies the row in the alert table that corresponds to this notification into the command. Notification Persistence

To prevent duplication and missed notifications after restart or failover, you must configure the Data Server for alert persistence in the **ALERTS** tab of the RTView Configuration Application.

Troubleshooting

This section includes:

- "Log Files for Solace," next
- "JAVA_HOME" on page 38
- "Permissions" on page 38
- "Network/DNS" on page 38
- "Data Not Received from Data Server" on page 39
- "Stop the Monitor" on page 39

Log Files for Solace

When any RTView Monitor for Solace component encounters an error, an error message is output to the console and/or to the corresponding log file. Logging is enabled by default. If you encounter issues with log files, verify the **logs** directory exists.

Solace Monitor Log Files

If you encounter issues, look for errors in the following log files, located in the **RTViewSolaceMonitor/projects/rtview-server/logs** directory:

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

RTView Manager Log Files

If you encounter issues, look for errors in the following log files, located in the **RTViewSolaceMonitor/projects/rtvmgr-server/logs** directory:

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

JAVA_HOME

If you encounter issues starting Solace Monitor or RTView Manager processes on Linux, verify that JAVA_HOME is set correctly in the path as JAVA_HOME is required for Tomcat to start correctly. On Windows, JAVA_HOME or JRE_HOME should exist as environment variables indicating a valid Java path.

Permissions

If you encounter permissions-related errors in the response from the **start_servers** 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 also verify confirm with your network administrator that you're not being blocked from accessing the remote system.

Data Not Received from Data Server

In the RTView Monitor for Solace, if you go to the **Administration>**"RTView Cache Tables" display and see that caches are not being populated with monitoring data (the number of rows in the table is zero), check for connection property errors that are provided to the Data Server:

1. "Open the RTView Monitor for Solace Project" and select **Solace Message Router** in the navigation tree. The **CONNECTIONS** tab opens.
2. Verify the connection parameters associated with your message routers.
3. Verify the SEMP version is correct for each of your message routers (monitoring data cannot be collected if the SEMP version is incorrect).
4. Click

in the title bar when finished.
5. Click

to apply your settings. It takes about 10-15 seconds for the data server to be available again.
6. In the RTView Manager for Solace (http://localhost:8068/rtview/solmon_manager), return to the **Administration>**"RTView Cache Tables" display and verify that all caches are being populated with monitoring data (the number of rows in the table is greater than zero).

Stop the Monitor

These instructions describe how to stop the RTView Monitor for Solace, RTView Manager and Tomcat by executing one command.

To stop the Monitor and Tomcat:

1. "Initialize a Command Prompt or Terminal Window".
2. Change directory (**cd**) to the **RTViewSolaceMonitor/bin** directory.
3. Execute **stop_servers.sh** (or **stop_servers.bat** for Windows) to stop all Monitor components, RTView Manager and Tomcat.
4. Optionally, you can use **grep** or **Task Manager** to ensure that all RTView-related processes and Tomcat are stopped.

- **UNIX**: Execute **ps –ef |grep rtv** to determine the Process Identifier of the processes still running and **kill -9 <ProcessId>** to terminate any that remain active.
- **Windows**: Open Task Manager and look for Java sessions with **hsqldb** or **rtv** in the execute statement and terminate any that remain active.

Start the Monitor

These instructions describe how to start the RTView Monitor for Solace (for tracking the health of your Solace resources) and the RTView Manager (for tracking the health of RTView Solace Monitor processes and Tomcat) using the pre-configured settings.

You execute one command to start the RTView Monitor for Solace, RTView Manager and Tomcat (the Monitor servlets are pre-deployed in Tomcat).

To start the Monitor and Tomcat:

1. "Initialize a Command Prompt or Terminal Window".

2. Change directory (**cd**) to the **RTViewSolaceMonitor/bin** directory.

3. Execute **sh start_servers.sh** (or **start_servers.bat** for Windows) to start all Monitor components, RTView Manager and Tomcat.

Important: UNIX/Linux - To make the script in the **bin** directory executable, use the **sh** command (as shown), or execute **chmod a+x start_servers.sh**, then execute **./ start_servers.sh**.

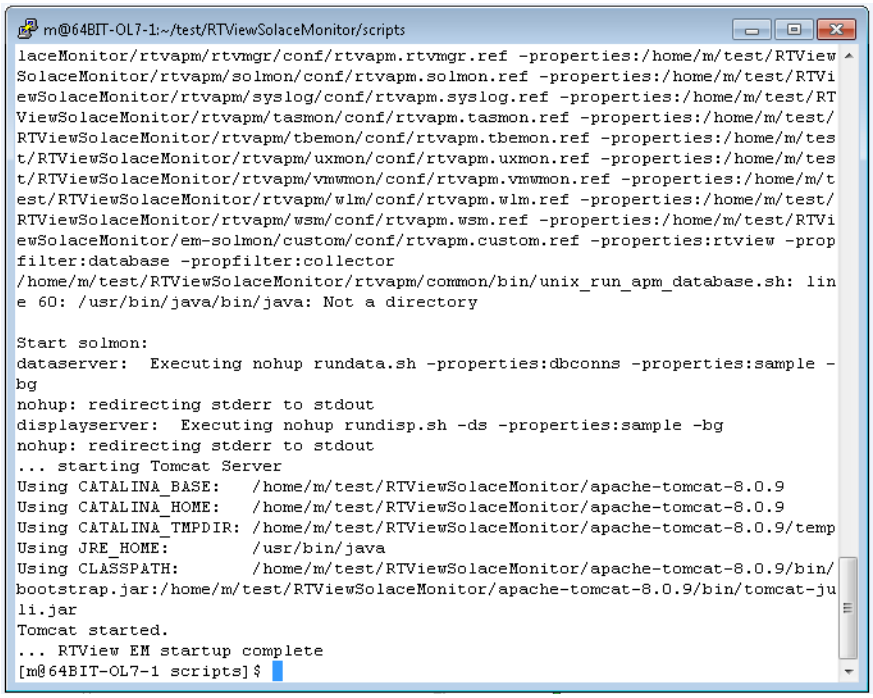

4. Open a browser and go to **localhost:8068/rtview/solmon_manager** (login ID/ Password is **solmon**/**solmonpw**). The Solace Monitor opens.

5. Go to the **Administration>**"RTView Cache Tables" display and verify that all caches are populated with monitoring data (the number of rows in the table is greater than zero). If not, there is a problem with the connection to the Data Server and/or the connection properties you created.

6. Open another browser window and go to **localhost:8068/rtview/rtvmgr_manager** (login ID/Password is **solmon**/**solmonpw**). The RTView Monitor opens.

7. Verify that all caches are populated with data.

CHAPTER 4 Additional Configurations

This section contains the following:

- "Obtain SEMP Version"
- "Create Instance from RTView Monitor for Solace"

Obtain SEMP Version

Skip this step if your Solace message routers are using Solace VMR version 8.7+ and Solace Appliance version 8.3+. This step is required if your Solace message routers are using software versions prior to Solace VMR version 8.7+ and Solace Appliance version 8.3+

Note: It is recommended to not include a SEMP version string in commands, and only include one if a known deprecated behavior is needed from a particular SEMP schema.

In order to properly request monitored data, the Monitor requires the exact SEMP version on your message routers. These instructions describe how to use SolAdmin to determine the SEMP version for each of your Solace Message Routers or VMRs. You will need this information when you connect your message routers and edit connection properties.

Note: These instructions are for SolAdmin on Windows. For Linux, only the path to the log file changes.

1. Navigate to the SolAdmin installation folder. For example, **C:\Program Files (x86)\SolAdmin**.

2. Change directory (**cd**) to the **bin** directory and open the **log4j.properties** file in a text editor.

3. Change the logging level to **DEBUG** and provide the full path to the logging file (for example, **C:\Logs**) while retaining all other settings. The edited properties are as follows:

full path to the location where you want the log file to be stored. In this example C:\Logs

log4j.appender.A1.File=**C:\Logs\soladmin.log**
# Set the logging category to DEBUG
log4j.category.com.solacesystems=**DEBUG**, A1

4. Save the **log4j.properties** file.

5. Start SolAdmin and add your message routers or VMRs as managed instances.

6. Open the **soladmin.log** file and locate the semp-version tag in SEMP requests. The SEMP version that will be used by the Monitor replaces underscores (**_**) with dots (**.**). For

example, if the SEMP request in the SolAdmin log file is **7_2VMR**, you use **7.2VMR** for the **$solSempVersion** substitution of the Monitor connection property.

Return to "Quick Start".

Create Instance from RTView Monitor for Solace

This section describes how to create obtain the RTView Monitor for Solace Amazon Machine Image (AMI).

On-premise users can skip this step.

The RTView Monitor for Solace AMI is pre-installed on an Amazon EC2 Amazon Machine Image (AMI) running Amazon Linux. It includes the following application stack for convenience of quick deployment:

- Oracle Java 8
- Node.js
- Docker
- MySQL 5.7 (via Docker) for storage of historical data
- rtvHostAgent (via Docker) for providing host metrics to RTVMGR
- cadvisor-rtview (via Docker) for providing docker metrics to RTVMGR

RTView Monitor for Solace AMI is configured to start all RTView processes and supporting services on restart.

The scripts used to create the Docker containers are included in named subdirectories under **/home/ec2-user/amibase**, to be used as templates if you wish to recreate the containers with your preferred configuration.

The MySQL database data is stored external to the Docker container at **/home/ec2-user/amibase/mysql/DATA**.

Before you proceed: We recommended that you be logged into your Amazon AWS user account with administrative access before following the link to the AWS Instance Launch Wizard.

1. In a browser, go to **http://sl.com/solace-ami-free-trial/** and complete the form to gain access to the page of region links.

2. Click on the link for the AWS region appropriate for you to go to the AWS Instance Launch Wizard.

3. In the **Configure Instance Details** screen, choose an appropriate **Instance Type** and click **Next: Configure Instance Details**.

For information about Instance Types, refer to AWS documentation. We recommend starting with the t2 family, of at least t2.medium.

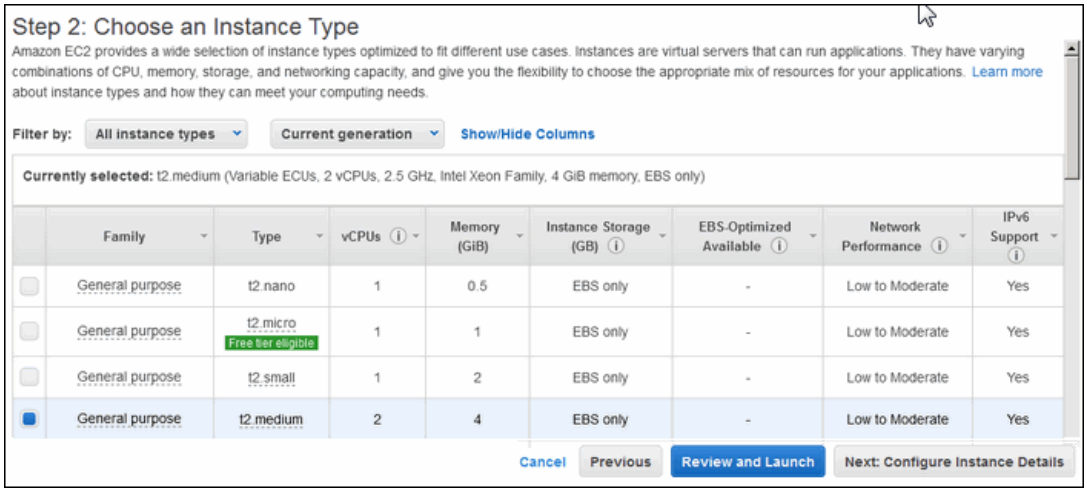

2. In the **Configure Instance Details** screen, configure the VPC, then click **Next: Add Storage**.

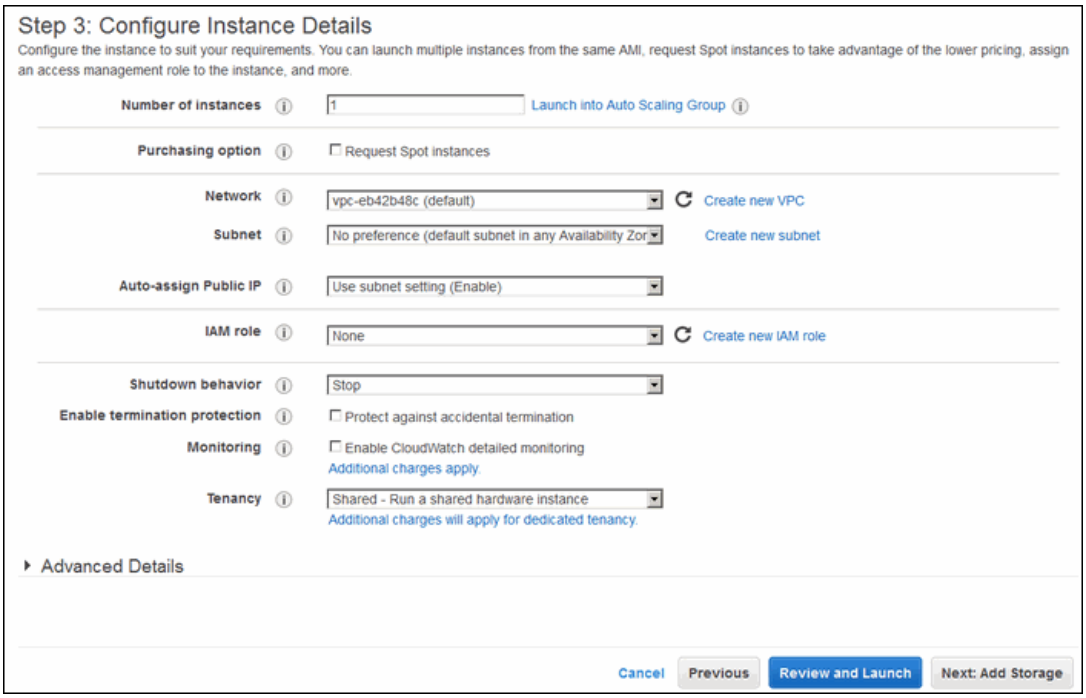

4. In the **Add Storage** screen, accept the **8 GB** storage size, or select a sufficiently-sized volume for the number of Solace message routers that you will be storing archival data for, and then click **Next: Tag Instance**.

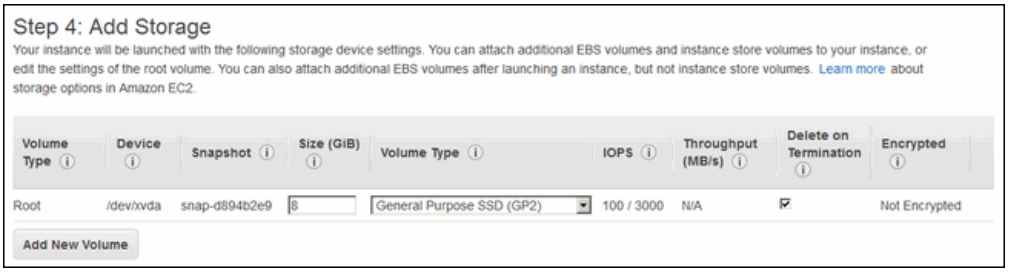

5. In the **Tag Instance** screen, add tags as appropriate to keep your VMR instances organized, then click **Next: Configure Security Group**.

   The following example uses Name, and Version but you can choose any tags that make sense for your application.

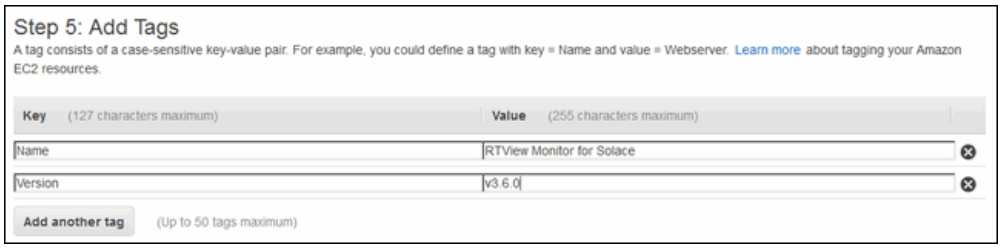

6. In the **Configure Security Group** screen, create or choose an appropriate security rule that allows SSH (22) and HTTP (80) access for the RTView Monitor for Solace, then click **Review and Launch**.

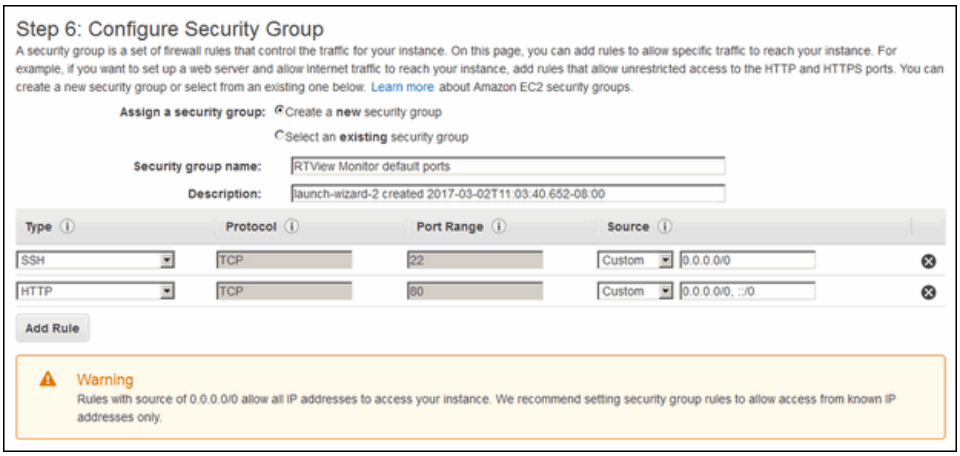

7. In the **Review** screen, verify your instance, ignore the warnings, and click **Launch**.

   The instance starts.

8. In the dialog box that opens, choose an authentication key pair for the instance, which can be used for this first login to the instance, then click **Launch Instance**.

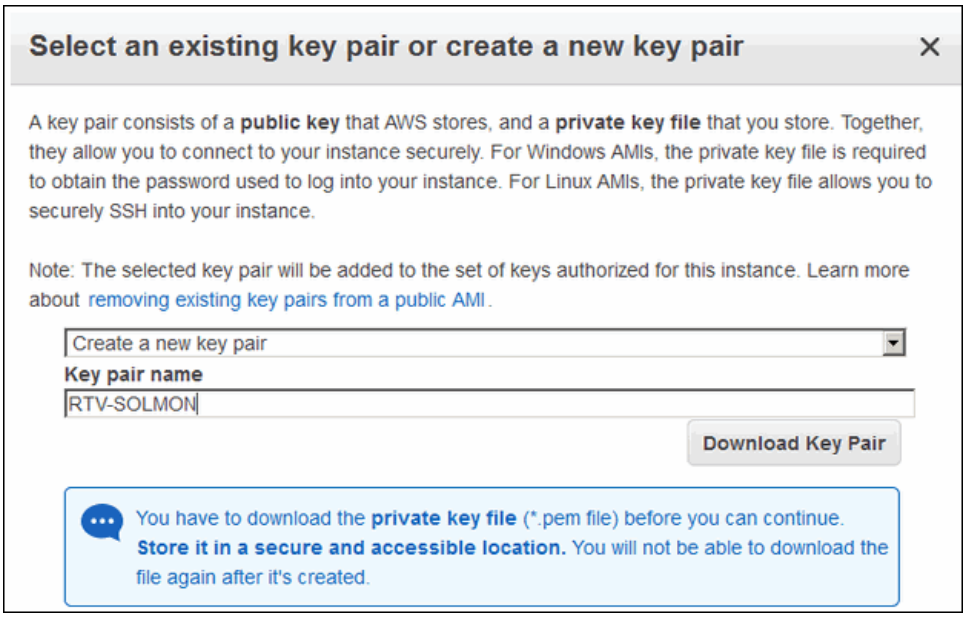

9. Look for your RTView Monitor for Solace instance in the EC2 dashboard **Instances**. This is where you can see the external and internal IP address of the instance.

For more information about IP Addressing in the Cloud, refer to Solace Corporation documentation.

10. To log into the Linux Host shell, enter the following command:

ssh -p 22 -i <auth_key> ec2-user@<public_ip>

Return to Quick Start instructions, "Add Message Router Connections".

CHAPTER 5 Using the Monitor

The RTView® Monitor for Solace® is an advanced messaging platform that allows customer applications to efficiently exchange messages over dedicated VPNs. The RTView® Monitor for Solace® provides pre-configured alerts and dashboards to monitor current status and manage history for the Solace message router. The RTView® Monitor for Solace® can help operators avoid or detect many problems relating to configuration, topology, and performance. This section describes Monitor features, graphs and functionality as well as Monitor displays.

This section includes:

- "Overview": This section describes the Monitor and GUI elements.
- "RTView Monitor for Solace Views/Displays": This section describes displays that are used by monitoring teams to monitor the health of Solace components (message routers, bridges, clients, endpoints and VPNs). To access RTView® Monitor for Solace®:
  - **http://ip_address:8068/rtview/solmon** if you are running the monitor remotely.
  - **http://localhost:8068/rtview/solmon** if you are running the monitor locally.

    Use **solmon**/**solmonpw** for username/password.

  Note: If you are using the RTView Monitor for Solace AMI version, you can also monitor "MySQL Database" and "Docker Engines" displays. If you are using the On-premise version these displays have no data.
- "RTView Manager for Solace Displays": This section describes displays that are used by administrators to set alert thresholds for RTView® Monitor for Solace®, including Syslog. To access RTView Manager for Solace:
  - **http://ip_address:8068/rtview/solmon_manager** if you are running the monitor remotely.
  - **http://localhost:8068/rtview/solmon_manager** if you are running the monitor locally.

    Use **solmon**/**solmonpw** for username/password.
- "RTView Manager Views/Displays": This section describes displays that are used by administrators to monitor the health of RTView® Monitor for Solace®. That is, to monitor components of the monitoring system itself (RTView servers, JVMs, Tomcat servers, hosts, Docker, MySQL and alert settings for these components). To access RTView Manager:
  - **http://ip_address:8068/rtview/rtvmgr_manager** if you are running the monitor remotely.
  - **http://localhost:8068/rtview/rtvmgr_manager** if you are running the monitor locally.

    Use **solmon**/**solmonpw** for username/password.

Overview

This section describes the general operation of the Solace Monitor and the user interface. This section includes:

- "Heatmaps": Describes how to read and use heatmaps.
- "Tables": Describes how to read and use tables.
- "Trend Graphs": Describes how to read and use trend graphs.
- "GUI Icons and Buttons": Describes title bar features and GUI elements shared by Monitor displays.

Heatmaps

Heatmaps organize your Solace resources (instances, databases, and collections) into rectangles and use color to highlight the most critical value in each. Heatmaps enable you to view various alert metrics in the same heatmap using drop-down menus. Each metric has a color gradient bar that maps relative values to colors. In most heatmaps, the rectangle size represents the number of resources in the rectangle; a larger size is a larger value. Heatmaps include drop-down menus by which to filter data. The filtering options vary among heatmaps (the **Solace Message Routers Heatmap** is shown below).

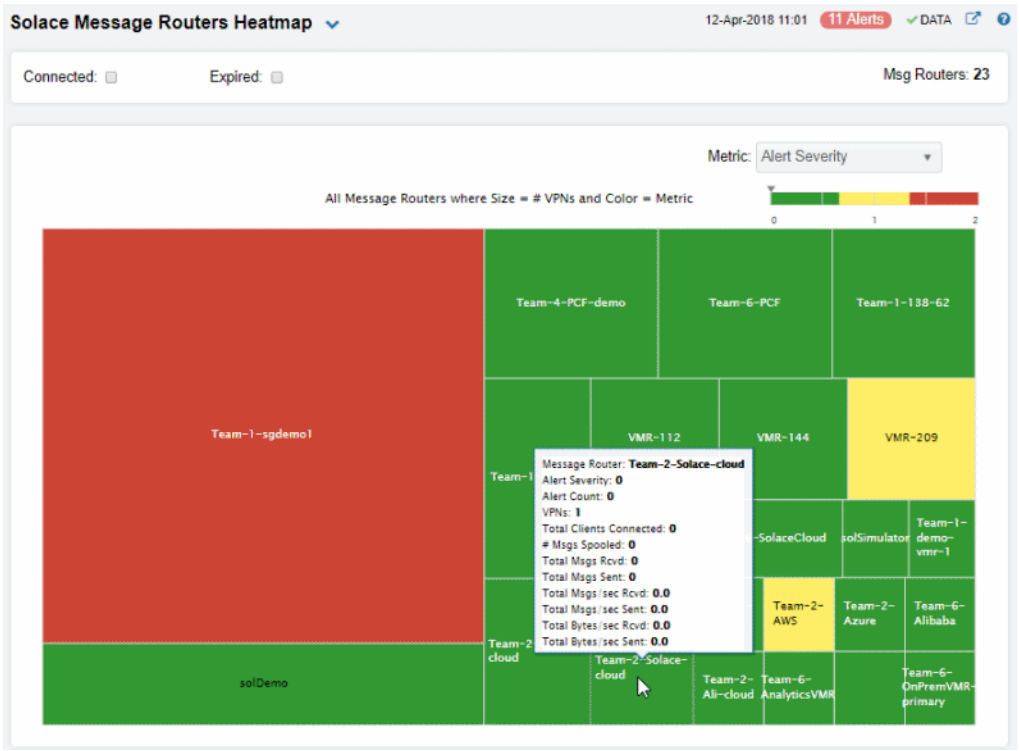

For example, the **Solace Message Routers Heatmap** contains a **Metric** drop-down menu with options such as **Alert Severity** and **Alert Count**. Menu options vary according to the data populating the heatmap. **Alert Severity** is selected and its corresponding color gradient 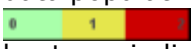 bar is shown. Each rectangle represents a connection. A red rectangle in the heatmap indicates that one or more resources associated with that connection currently has an alert in an alarm state. The yellow rectangles in the heatmap indicate that one or more resources associated with that host currently have an alert in a warning state. A green rectangle would indicate that no alert is in a warning or alarm state.

In most heatmaps, you can also drill-down to more detail by clicking a rectangle in the heatmap.

Note: Typically, it takes about 30 seconds after a server is started to appear in an Solace Monitor display. By default, data is collected every 15 seconds, and the display is refreshed 15 seconds afterward.

As previously mentioned, each Metric drop-down menu option has a color gradient bar that maps relative values to colors. The following summarizes the heatmap color code translation for typical heatmaps:

Alert Impact

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

Alert Severity

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

Metrics that have exceeded their specified **ALARM LEVEL** threshold have an **Alert Severity** value of **2**. For a given rectangle, this indicates that one or more metrics have reached their alert thresholds.

Metrics that have exceeded their specified **WARNING LEVEL** threshold have an **Alert Severity** value of **1**. For a given rectangle, this indicates that one or more metrics have reached their warning thresholds.

Metrics that have not exceeded either specified threshold have an **Alert Severity** value of **0**. For a given rectangle, this indicates that no metrics have reached their warning or alert thresholds.

Alert Count

The total number of critical and warning alerts in a given item (index) associated with the rectangle. The color gradient bar numerical values range from **0** to the maximum count of alerts currently in the heatmap. The middle value in the gradient bar indicates the average alert count.

Criticality

The maximum level of **Criticality** (rank of importance) in a given item (index) associated with the rectangle. Values range from **0** to **5**, as indicated in the color gradient bar, where **5** is the highest Criticality.

Criticality is specified in the Service Data Model by your administrator. **Criticality** values range from **A** to **E**, where **A** is the highest Criticality (level **5** maps to a Criticality of **A** and level **1** maps to a **Criticality** of **E** with equally spaced intermediate values).

Mouse-over

The mouse-over functionality provides additional detailed data in a tool-tip when you mouse-over a heatmap. The following figure illustrates mouse-over functionality in a heatmap object. In this example, when you mouse-over a host, details are shown such as alert count, number of connections, and pending messages.

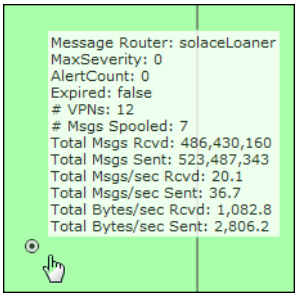

Tables

Solace Monitor tables contain the same data that is shown in the heatmap in the same View, and additional data not included in the heatmap. For example, the **VPNs Table** display (shown below) shows the same data as the **VPNs Heatmap** display.

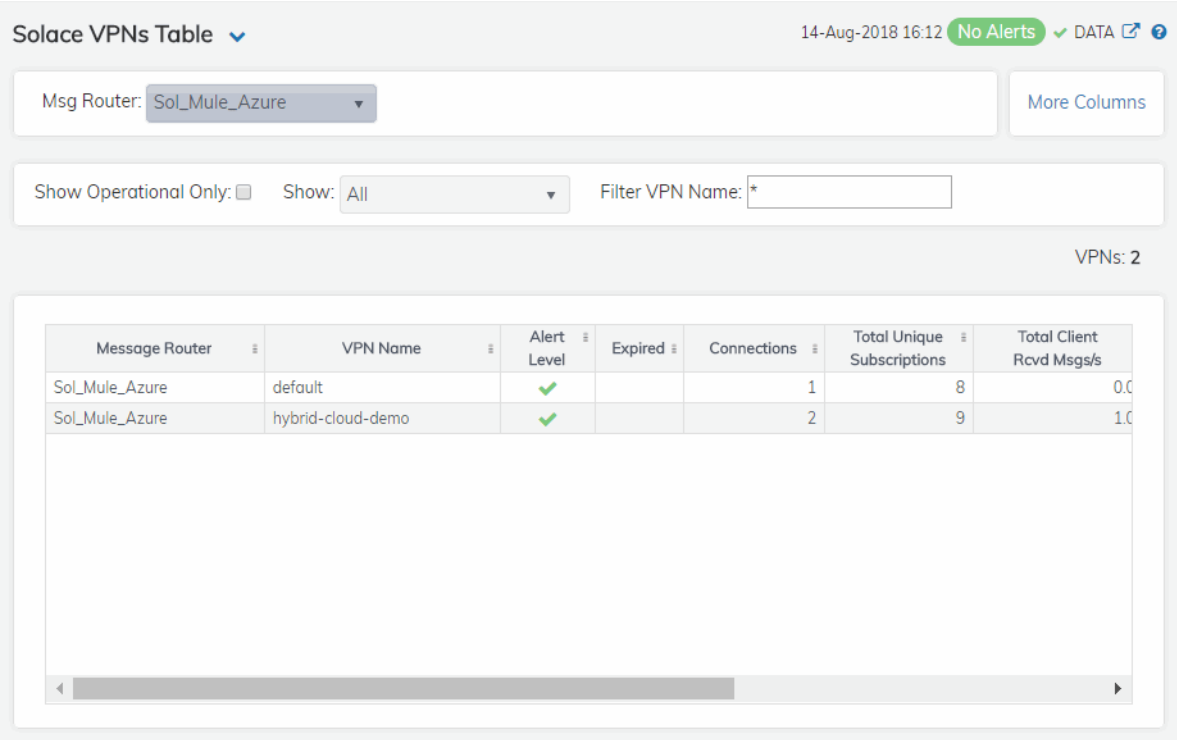

Tables support advanced HTML, interactive features: sorting on multiple columns, filtering on multiple columns, column resizing, column reordering, and hiding columns. Many of these features are accessed from the column menu, shown in the screen shot above, which you open by clicking on the menu icon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

Additional features are:

- "Multiple Column Sorting," next
- "Column Visibility" on page 51
- "Column Filtering" on page 51
- "Column Locking" on page 53
- "Column Reordering" on page 53
- "Saving Settings" on page 54
- "Row Paging" on page 54

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.

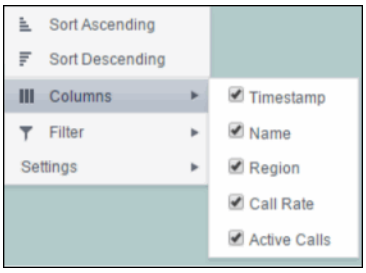

The leftmost column (the row header column) cannot be hidden.

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:

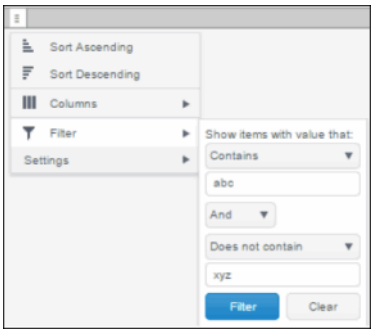

Options in the **Column Filter** dialog vary according to the data type of the selected column:

- **String columns**: You can enter a filter string such as "abc" and, from the drop-down 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.

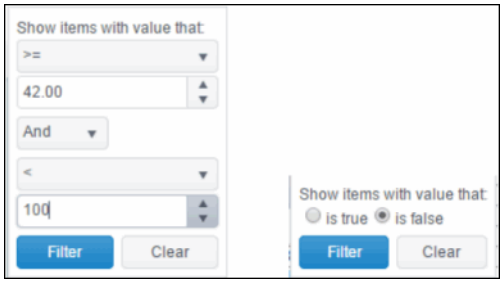

- **Date columns**: You can select a date and time and choose whether matching items should have a timestamp that is the same as, before, or after the filter time. The date is selected by clicking on the calendar icon and picking a date from a calendar dialog. The time is selected by clicking on the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

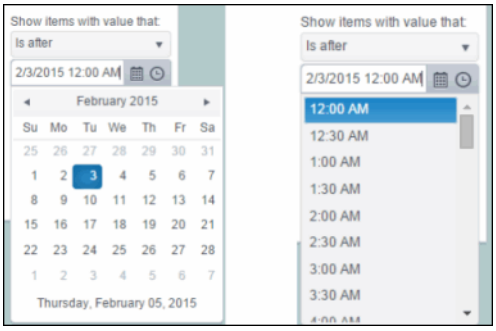

  Alternatively, a date and time can be typed into the edit box. The strings shown in a date column are formatted by the Display Server using its time zone. But if a filter is specified on a date column, the date and time for the filter are computed using the client system's time zone. This can be confusing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

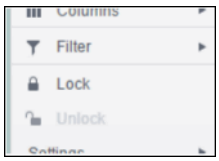

If the row header is enabled, at least one column must remain locked.

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**:

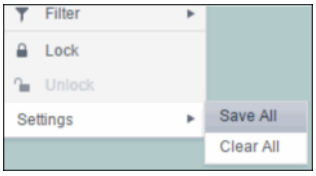

The grid's settings are written as an item in the browser's local storage. The item's value is a string containing the grid's settings. The item uses a unique key comprised of the URL path name, the display name, and the table's RTView object name. If the Thin Client's login feature is enabled, the key 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.

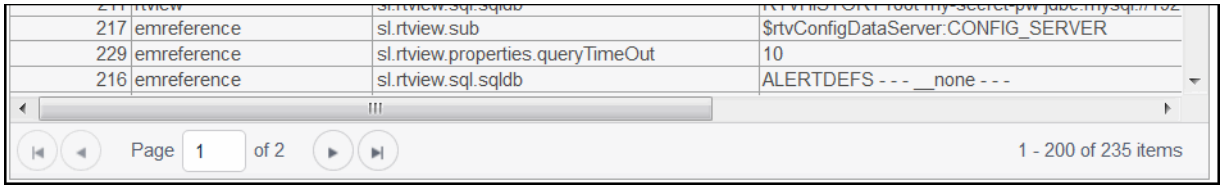

Trend Graphs

Solace Monitor trend graphs enable you to view and compare various important metrics over time, such as server memory and virtual memory utilization.

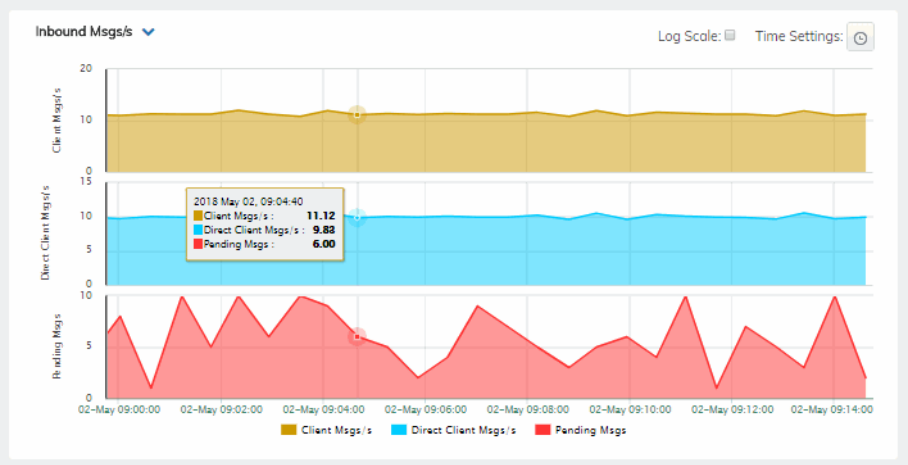

Time Settings

By default, the time range end point is the current time. To change the time range, click the **Time Settings** and either:

- choose a **Time range** from 5 Minutes to 7 Days in the drop-down menu.
- specify begin/end dates using the calendar ..
- specify begin/end time using the clock .

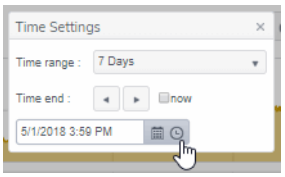

Toggle forward/backward in the trend graph per the period you choose (from the **Time range** drop-down menu) using arrows .

Restore settings to current time by selecting **now** .

Mouse-over

The mouse-over functionality provides additional detailed data in an over imposed pop-up window when you mouse-over trend graphs.

Log Scale

The Log Scale option enables visualization on a logarithmic scale. This option should be used when the range in your data is very broad. For example, if you have data that ranges from the tens to the thousands, then data in the range of tens will be neglected visually if you do not check this option. This option makes data on both extreme ranges visible by using the logarithmic of the values rather than the actual values.

GUI Icons and Buttons

Displays share the same title bar as shown below.

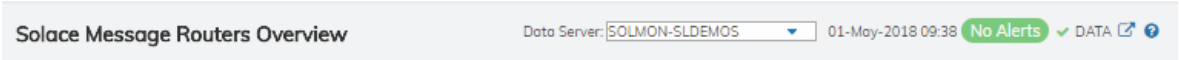

The following describes GUI icons and behavior in the title bar.

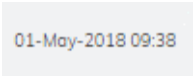

The current local date and time. If the time is incorrect, this might indicate that the monitor stopped running. When the date and time is correct and the **Data** indicator is green, this is a strong indication that the platform is receiving current and valid data.

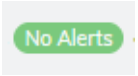

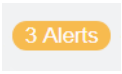

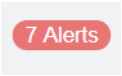

ALERTS: Opens the Alerts Table, shows the total number of alerts associated with items currently in the display as well as the maximum alert severity of these, where:

Green indicates that no metrics have exceeded their alert thresholds.

Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.

Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.

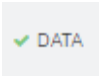

DATA: The data source is currently connected. When the date and time is correct and the **DATA** indicator is green, this is a strong indication that the platform is receiving current and valid data.

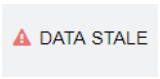

DATA STALE: The data source is currently disconnected. There has been no response from the Data Server for 31+ seconds.

This feature is currently in development.

This feature is currently in development.

Time Settings

By default, the time range end point is the current time. To change the time range, click the **Time Settings** and either:

- choose a **Time range** from 5 Minutes to 7 Days in the drop-down menu.
- specify begin/end dates using the calendar ..
- specify begin/end time using the clock .

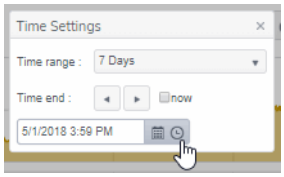

Toggle forward/backward in the trend graph per the period you choose (from the **Time range** drop-down menu) using arrows .

Restore settings to current time by selecting **now** .

Log Scale

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

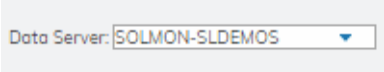

Drop-down menus for selecting the item/s you want to view. Options differ among displays.

RTView Monitor for Solace Views/Displays

The RTView® Monitor for Solace® home page provides a health summary of all your Solace message routers, as shown in the following figure.

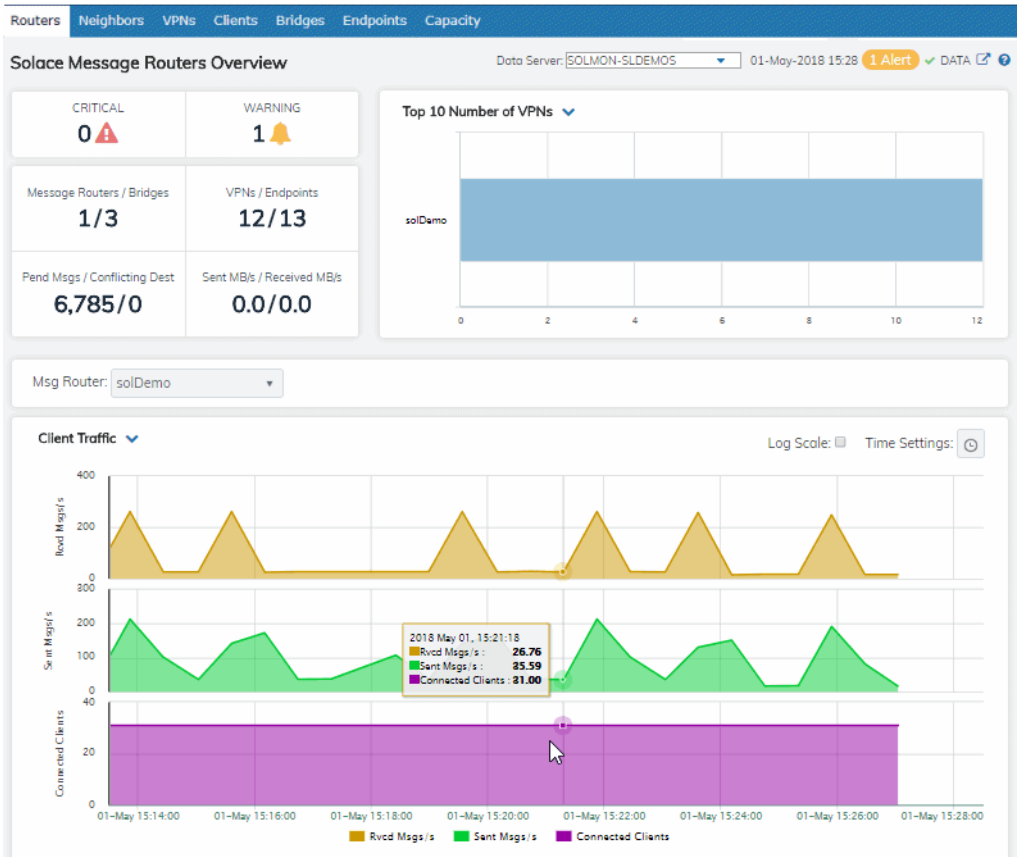

The RTView® Monitor for Solace® has the following Views:

- "Routers": The displays in this View present message router-level metrics, which reflect configuration settings, total throughput, current status, errors, and value-added calculations that summarize metrics across all of the VPNs.
- "Neighbors": The displays in this View present a topology and metrics of your message routers, VMRs and servers as well as and their configuration settings.
- "VPNs": The displays in this View present VPN-level metrics.
- "Clients": The displays in this View present metrics for all clients of the message router. These views can be filtered to limit the displays to clients for a single VPN.
- "Bridges": The displays in this View present a topology and metrics of your bridges and VPNs. These views can be filtered to limit the displays to bridges for a single VPN.
- "Endpoints": The displays in this View present metrics for topics and queues on the message router, which can be filtered to limit the displays to topics and queues for a single VPN.
- "Capacity": The displays in this View present current metrics, alert count and severity at the message router level.

Routers

These displays provide detailed metrics for message routers and their connected message routers. Displays in this View are:

- "Routers - Overview": Health snapshot of top 10 most utilized VPNs, trend graphs trace key performance metrics such as messages sent/received and connected clients.
- "Routers - Heatmap": A color-coded heatmap view of the current status of each of your message routers.
- "Routers - Table": A tabular view of all available message router performance data.
- "Routers - Summary": Current and historical metrics for a single message router.
- "Routers - Sensors": Provides value and status information for all sensors on a single message router or for all sensors for all message routers.
- "Routers - Provisioning": Provides message router details such as host, chassis, redundancy, memory, and fabric data for a particular message router.
- "Routers - Interface": Provides detailed data and status information for the interfaces associated with one or all message router(s). You can also view current and historical amounts of incoming and outgoing packets and bytes for a selected interface in a trend graph.
- "Routers - Message Spool": Provides status and usage data for message spools associated with one or all message router(s).

Routers - Overview

View a health snapshot of top 10 most utilized VPNs, trend graph traces key performance metrics such as messages sent/received and connected clients.

Select a data server, message router and metric from the drop-down menus. Consider keeping this display open for monitoring at a glance.

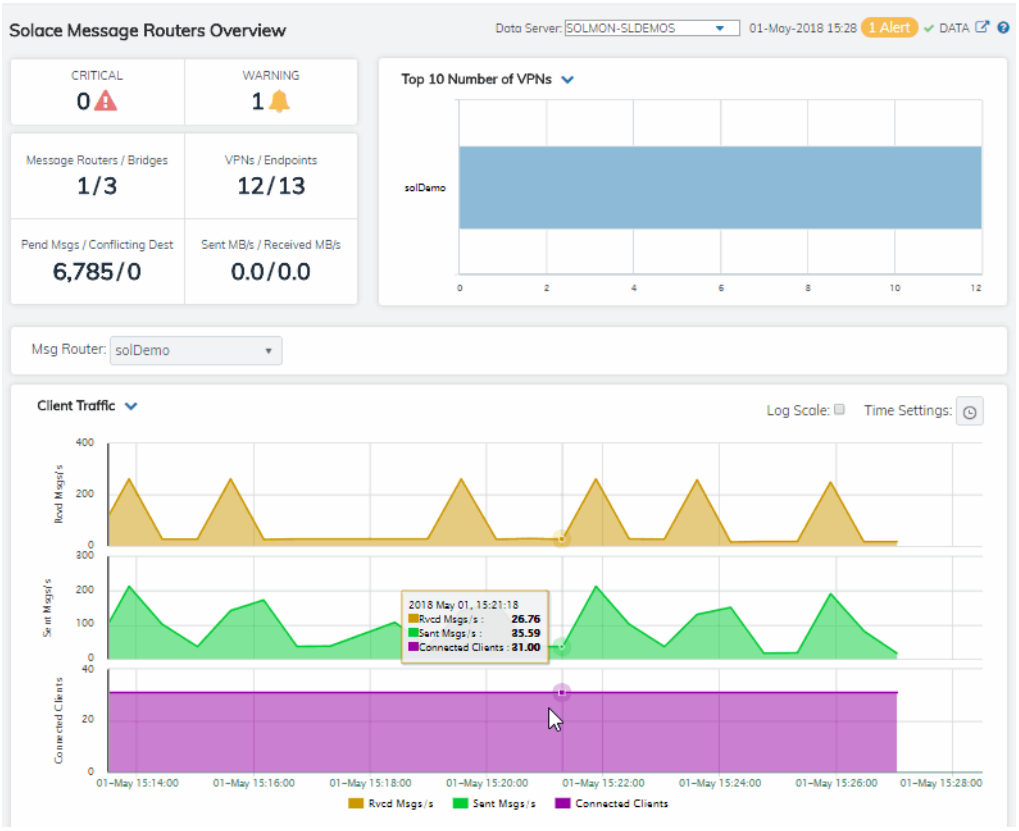

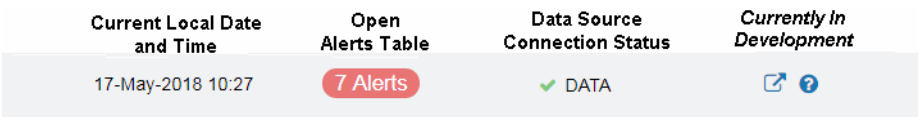

CRITICAL	Total number of current critical alerts for message routers on the selected data server.
WARNING	Total number of current critical alerts for message routers on the selected data server.
Message Routers/Bridges	Total number of message routers/bridges on the selected data server.
VPNs/Endpoints	Total number of VPNs/endpoints on the selected data server.
Pending Msgs/Conflicting Dest	Total number of pending messages/conflicting destinations on the selected data server.
Sent MBs/Received MBs	Total number of MBs sent/MBs received on the selected data server.
Top 10 Number of VPNs	Ten message routers with the greatest number of connected VPNs.

Msg Router

Select a message router to trace performance metrics in the trend graph, then choose a metric:

Client Traffic: Traces the number of messages received per second, messages sent per second and the number of connected clients.

Spool Msgs: Traces the number of spooled messages and spool size (in megabytes.)

Time Settings

By default, the time range end point is the current time. To change the time range, click the **Time Settings** and either:

- choose a **Time range** from 5 Minutes to 7 Days in the drop-down menu.
- specify begin/end dates using the calendar .
- specify begin/end time using the clock .

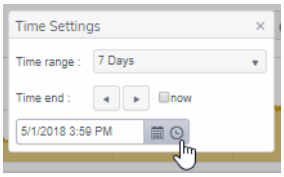

Toggle forward/backward in the trend graph per the period you choose (from the **Time range** drop-down menu) using arrows .

Restore settings to current time by selecting **now** .

Log Scale

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

Routers - Heatmap

View the current status of all your message routers. Each rectangle in the heatmap is a single message router where the rectangle size represents the number of its connected clients. The rectangle color reflects where the current value is on its color gradient bar. Select a router from the drop-down menu. For example, by default, **Alert Severity** is shown:

Alert Severity

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

Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.

Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.

Green indicates that no metrics have exceeded their alert thresholds.

Each metric has its own color gradient bar (scroll down for more "Metric Options").

Mouse over a rectangle to see additional metrics. Use the check-boxes ☑ to include / exclude **Connected** and **Expired** message routers. Click a rectangle to drill-down to details about a message router in the "Routers - Summary" display.

Consider keeping this display open for monitoring at a glance.

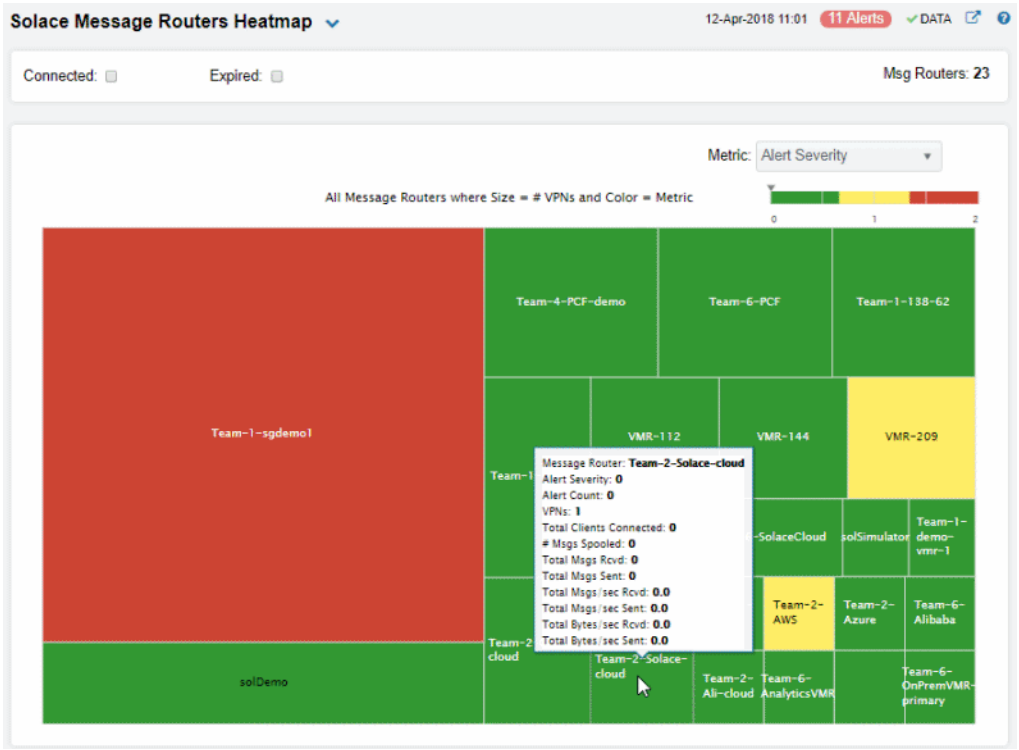

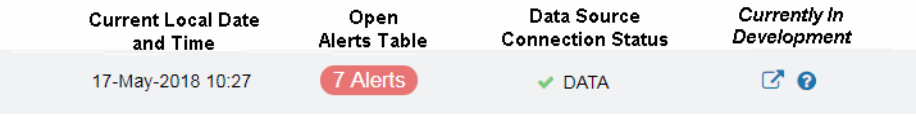

Metric Options

Choose a metric from the drop-down menu:

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

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

Alert Count The total number of critical and warning alerts. The color gradient 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.

 # Msgs Spooled The total number of spooled messages. The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the defined alert threshold of **SolMsgRouterPendingMsgsHigh**. The middle value in the gradient bar indicates the middle value of the range.

Total Msgs Rcvd The total number of received messages. The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of total messages received in the heatmap. The middle value in the gradient bar indicates the average count.

Total Msgs Sent	The total number of sent messages. The color gradient 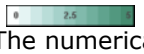 bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of total messages sent in the heatmap. The middle value in the gradient bar indicates the average count.
Total Msgs/ sec Rcvd	The number of messages received per second. The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the defined alert threshold of **SolMsgRouterInboundMsgRateHigh**. The middle value in the gradient bar indicates the middle value of the range.
Total Msgs/ sec Sent	The total number of messages sent per second. The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the defined alert threshold of **SolMsgRouterOutboundMsgRateHigh**. The middle value in the gradient bar indicates the middle value of the range.
Total Bytes/ sec Rcvd	The total number of bytes received per second in the message router. The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **SolMsgRouterInboundByteRateHigh**. The middle value in the gradient bar indicates the middle value of the range.
Total Bytes/ sec Sent	The total number of bytes sent per second in the message router. The color gradient 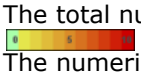 bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **SolMsgRouterOutboundByteRateHigh**. The middle value in the gradient bar indicates the middle value of the range.

Routers - Table

View current status data for all message routers in a tabular format. Each row in the table is a different message router.

By default, a subset of available metrics is shown. Use 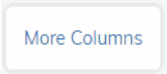

to toggle to the complete set of metrics available (and back to the subset).

Msg Routers: 23 (in the upper right portion) is the number of message routers in the display.

Select a router from the drop-down menu. Sort data in numerical or alphabetical order on column headers. Use the check-boxes ☑ to include / exclude **Connected** and **Expired** message routers.

Data shown in the "Routers - Heatmap" is shown as well as additional metrics.

Double-click a row to drill-down and investigate in the "Routers - Summary" display. See "Column Values" for details about metrics shown.

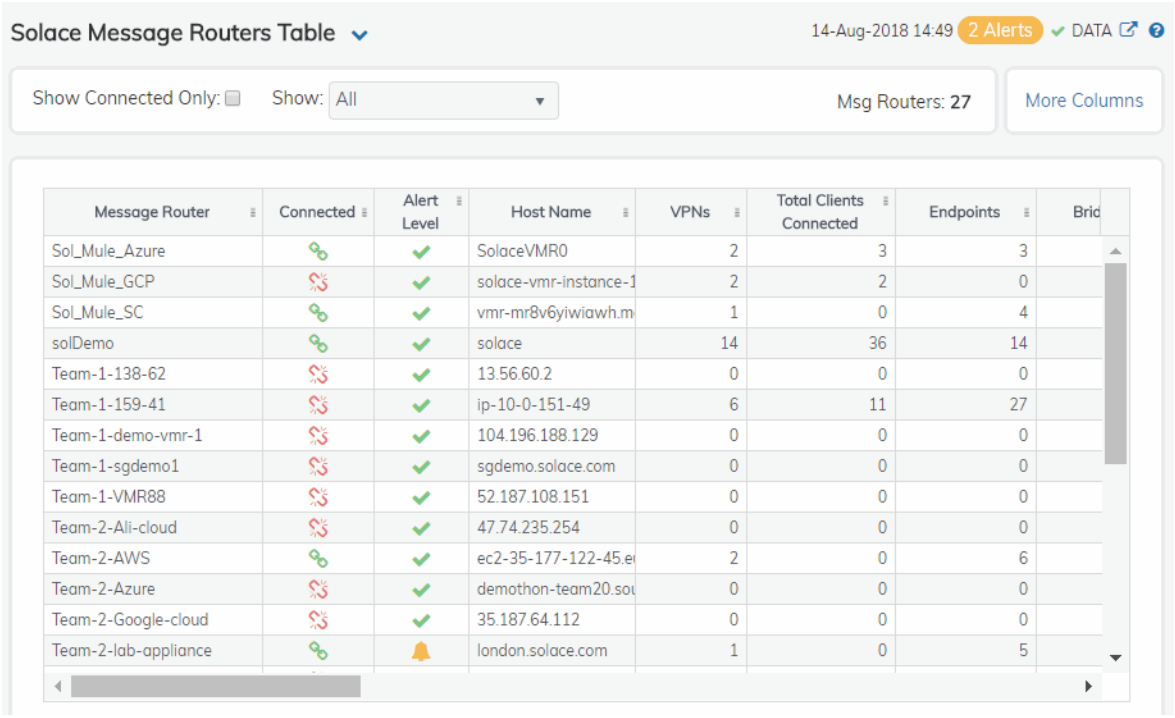

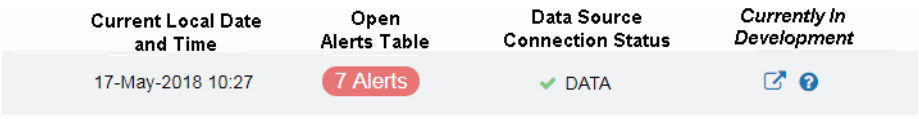

Column Values

Message Router The name of the message router.

Connected The message router state:
- Red indicates that the message router is NOT connected.
- Green indicates that the message router is connected.

Alert Severity The current alert severity:
- Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.
- Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.
- Green indicates that no metrics have exceeded their alert thresholds.

Alert Count The total number of alerts.

Expired When checked, performance data about the sensor has not been received within the time specified (in seconds) in the **$solRowExpirationTime** field in the **conf\rtvapm_solmon.properties** file. The **$solRowExpirationTimeForDelete** field allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response from the sensor. To view/edit the current values, modify the following lines in the **.properties** file:

```
# Metrics data are considered expired after this number
of seconds
#
collector.sl.rtview.sub=$solRowExpirationTime:45
collector.sl.rtview.sub=$solRowExpirationTimeForDelete
:3600
```

In the example above, the **Expired** check box would be checked after 45 seconds, and the row would be removed from the table after 3600 seconds.

Host Name The name of the host.

Platform The name of the platform.

OS Version The version of the operating system.

Up Time The amount of time that the message router has been up and running.

VPNs The total number of VPNs configured on the message router.

Total Clients The total number of clients associated with the message router.

Total Clients Connected The total number of clients that are currently connected to the message router.

Clients Using Compression The number of clients who send/receive compressed messages.

Clients Using SSL The number of clients using SSL for encrypted communications.

Max Client Connections The maximum number of available client connections.

Endpoints The total number of Endpoints configured on the message router.

Bridges The total number of bridges configured on the message router.

Local Bridges The total number of local bridges configured on the message router.

Remote Bridges	The total number of remote bridges configured on the message router.
Remote Bridge Subscriptions	The total number of remote bridge subscriptions configured on the message router.
Routing Enabled	This check box is checked when the message router is configured to route messages to other message routers.
Routing Interface	The name of the interface configured to support message routing.
Total # Conflicting Destinations	The total number conflicting destinations.
Pending Messages	The number of pending messages on the message router.
Total Client Msgs Rcvd	The total number of client messages received on the message router.
Total Client Msgs Sent	The total number of client messages sent by the message router.
Total Client Msgs Rcvd/sec	The total number of client messages received per second by the message router.
Total Client Msgs Sent/ sec	The total number of client messages sent by the message router.
Total Client Bytes Rcvd	The total number of client bytes received by the message router.
Total Client Bytes Sent	The total number of client bytes sent by the message router.
Total Client Bytes Rcvd/sec	The total number of client bytes received per second by the message router.
Total Client Bytes Sent/sec	The total number of client bytes sent per second by the message router.
Total Client Direct Msgs Rcvd	The total number of direct client messages received by the message router.
Total Client Direct Msgs Sent	The total number of direct client messages sent from the message router.
Total Client Direct Msgs Rcvd/sec	The total number of direct client messages received per second by the message router.
Total Client Direct Msgs Sent/sec	The total number of direct client messages sent per second by the message router.
Total Client Direct Bytes Rcvd	The total number of direct client bytes received by the message router.
Total Client Direct Bytes Sent	The total number of direct client bytes sent by the message router.
Total Client Direct Bytes Rcvd/sec	The total number of direct client bytes received per second by the message router.
Total Client Direct Bytes Sent/sec	The total number of direct client bytes sent per second by the message router.
Total Client Non-Persistent Msgs Rcvd	The total number of non-persistent client messages received by the message router.
Total Client Non-Persistent Msgs Sent	The total number of non-persistent client messages sent by the message router.
Total Client Non-Persistent Msgs Rcvd/sec	The total number of non-persistent client messages received per second by the message router.

Total Client Non-Persistent Msgs Sent/ sec	The total number of non-persistent client messages sent per second by the message router.
Total Client Non-Persistent Bytes Rcvd	The total number of non-persistent client bytes received by the message router.
Total Client Non-Persistent Bytes Sent	The total number of non-persistent client bytes sent by the message router.
Total Client Non-Persistent Bytes Rcvd/sec	The total number of non-persistent client bytes received per second by the message router.
Total Client Non-Persistent Bytes Sent/ sec	The total number of non-persistent client bytes sent per second by the message router.
Total Client Persistent Msgs Rcvd	The total number of persistent client messages received by the message router.
Total Client Persistent Msgs Sent	The total number of persistent client messages sent by the message router.
Total Client Persistent Msgs Rcvd/sec	The total number of persistent client messages received per second by the message router.
Total Client Persistent Msgs Sent/ sec	The total number of persistent client messages sent per second by the message router.
Total Client Persistent Bytes Rcvd	The total number of persistent client bytes received by the message router.
Total Client Persistent Bytes Sent	The total number of persistent client bytes sent by the message router.
Total Client Persistent Bytes Rcvd/sec	The total number of persistent client bytes received per second by the message router.
Total Client Persistent Bytes Sent/ sec	The total number of persistent client bytes sent per second by the message router.
Avg Egress Bytes/min	The average number of outgoing bytes per minute.
Avg Egress Compressed Msgs/min	The average number of outgoing compressed messages per minute.
Avg Egress Msgs/min	The average number of outgoing messages per minute.
Avg Egress SSL Msgs/min	The average number of outgoing messages per minute being sent via SSL-encrypted connections.
Avg Egress Uncompressed Msgs/min	The average number of uncompressed outgoing messages per minute.
Avg Ingress Bytes/min	The average number of incoming bytes per minute.
Avg Ingress Compressed Msgs/min	The average number of compressed incoming message per minute.
Avg Ingress Msgs/min	The average number of incoming messages per minute.
Average Ingress SSL Msgs/ min	The average number of incoming messages per minute being received via SSL-encrypted connections.
Avg Ingress Uncompressed Msgs/min	The average number of uncompressed messages per minute.
Current Egress Bytes/sec	The current number of outgoing bytes per second.

Current Egress Compressed Msgs/sec	The current number of outgoing compressed messages per second.
Current Egress Msgs/sec	The current number of outgoing messages per second.
Current Egress SSL Msgs/sec	The current number of outgoing messages per second sent via SSL-encrypted connections.
Current Egress Uncompressed Msgs/sec	The current number of outgoing uncompressed messages per second.
Current Ingress Bytes/sec	The current number of incoming bytes per second.
Current Ingress Compressed Msgs/sec	The current number of incoming compressed messages per second.
Current Ingress Msgs/sec	The current number of incoming messages per second.
Current Ingress SSL Msgs/sec	The current number of incoming messages per second received via SSL-encrypted connections.
Current Ingress Uncompressed Msgs/sec	The current number of incoming uncompressed messages per second.
Ingress Comp Ratio	The percentage of incoming messages that are compressed.
Egress Comp Ratio	The percentage of outgoing messages that are compressed.
Egress Compressed Bytes	The number of outgoing compressed bytes.
Egress SSL Bytes	The number of outgoing compressed bytes being sent via SSL-encrypted connections.
Egress Uncompressed Bytes	The number of outgoing uncompressed bytes.
Ingress Compressed Bytes	The number of incoming compressed bytes.
Ingress SSL Bytes	The number of incoming bytes via SSL-encrypted connections.
Ingress Uncompressed Bytes	The number of incoming uncompressed bytes.
Total Egress Discards	The total number of outgoing messages that have been discarded by the message router.
Total Egress Discards/sec	The total number of outgoing messages per second that have been discarded by the message router.
Total Ingress Discards	The total number of incoming messages that have been discarded by the message router.
Total Ingress Discards/sec	The total number of incoming messages per second that have been discarded by the message router.
Client Authorization Failures	The number of failed authorization attempts
Client Connect Failures (ACL)	The number of client connection failures caused because the client was not included in the defined access list.
Subscribe Topic Failures	The number of failed attempts at subscribing to topics.
TCP Fast Retrans Sent	The total number of messages that were retransmitted as a result of TCP Fast Retransmission (one or more messages in a sequence of messages that were not received by their intended party that were sent again).
Memory (KB)	The total available memory (in kilobytes) on the message router.

Memory Free (KB)	The total amount of available memory (in kilobytes) on the message router.
Memory Used (KB)	The total amount of memory used (in kilobytes) on the message router.
Memory Used %	The percentage of total available memory that is currently being used.
Swap (KB)	The total available swap (in kilobytes) on the message router.
Swap Free (KB)	The total amount of available swap (in kilobytes) on the message router.
Swap Used (KB)	The total amount of swap used (in kilobytes) on the message router.
Swap Used %	The percentage of total available swap that is currently being used.
Subscription Mem Total (KB)	The total amount of available memory (in kilobytes) that can be used by queue/topic subscriptions.
Subscription Mem Free (KB)	The current amount of available memory (in kilobytes) that can be used by queue/topic subscriptions.
Subscription Mem Used (KB)	The current amount of memory (in kilobytes) being used by queue/topic subscriptions.
Subscription Mem Used %	The percentage of available memory being used by queue/topic subscriptions.
Chassis Product Number	The product number of the chassis in which the router is contained.
Chassis Revision	The revision number of the chassis.
Chassis Serial	The serial number of the chassis.
BIOS Version	The basic input/output system used by the chassis.
CPU-1	The name of the central processing unit (CPU 1) used by the message router.
CPU-2	The name of the central processing unit (CPU 2) used by the message router.
Operational Power Supplies	The number of available power supplies that are operational on the chassis.
Power Redundancy Config	The configuration used by the backup message router.
Max # Bridges	The maximum number of bridges allowed on the message router.
Max # Local Bridges	The maximum number of local bridges allowed on the message router.
Max # Remote Bridges	The maximum number of remote bridges allowed on the message router.
Max # Remote Bridge Subscriptions	The maximum number of remote bridge subscriptions allowed on the message router.
Redundancy Config Status	The status of the redundancy configuration.
Redundancy Status	The status of the redundant message router.
Redundancy Mode	Refer to Solace documentation for more information.
Auto-revert	Refer to Solace documentation for more information.

Mate Router Name	If redundancy is configured, this field lists the redundant router name (mate router name).
ADB Link Up	This check box is checked if a message router is set up to use guaranteed messaging and an Assured Delivery Blade (ADB) is set up and working correctly.
ADB Hello Up	Refer to Solace documentation for more information.
Pair Primary Status	The primary status of the message router and its redundant (failover) mate.
Pair Backup Status	Refer to Solace documentation for more information.

Expired	When checked, performance data about the message router has not been received within the time specified (in seconds) in the **$solRowExpirationTime** field in the **conf\rtvapm_solmon.properties** file. The **$solRowExpirationTimeForDelete** field allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response from the message router. To view/edit the current values, modify the following lines in the **.properties** file:
	`# Metrics data are considered expired after this number of seconds`
	`#`
	`collector.sl.rtview.sub=$solRowExpirationTime:45`
	`collector.sl.rtview.sub=$solRowExpirationTimeForDelete:3600`
	
	In the example above, the **Expired** check box would be checked after 45 seconds, and the row would be removed from the table after 3600 seconds.

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

Routers - Summary

View trend graphs of current and historical performance metrics for client traffic. Check general health details for any message router.

Choose a message router from the drop-down menu to view total number of connected clients, number of incoming messages, **Up Time**, and additional information specific to a message router. You can also view alert statuses for the message router and **Spool Status** data for the message router.

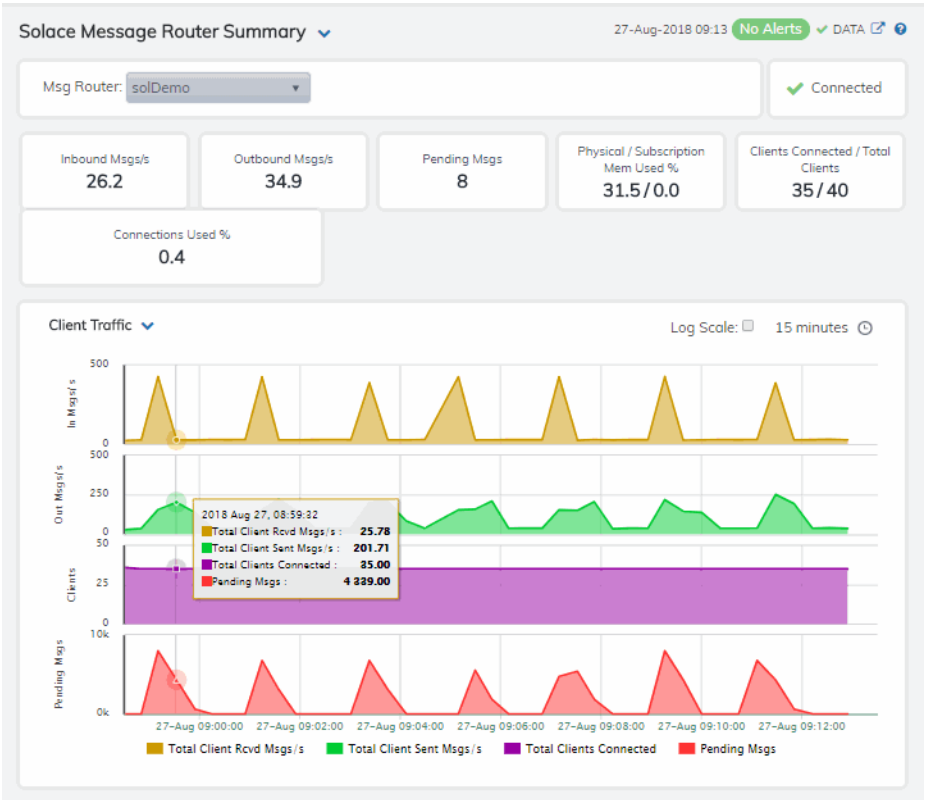

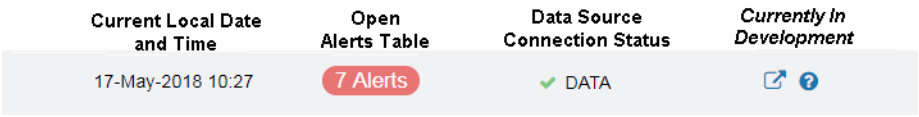

The connection status.

Inbound Msgs/s	The number of messages received per second.
Outbound Msgs/s	The number of messages sent per second.
Pending Msgs/s	The number of pending messages.
Physical / Subscription Mem Used %	The total percentage of physical memory used / the total percentage of subscription memory used.
Clients Connected / Total Clients	The current number of clients connected / the total number of clients.
Connections Used %	The percentage of connections used.

Trend Graphs
Traces the sum for the selected message router.

Client Traffic
- **In Msgs/s** - Traces the total number of client messages received per second.
- **Out Msgs/s** - Traces the total number of client messages sent per second.
- **Clients** - Traces the total number of connected clients.
- **Pending Msgs** - Traces the total number of pending messages.

Spool Msgs
- **Pending Msgs**- Traces the total number of pending spool messages.
- **Spool Usage MB** - Traces the total amount of space used by spool messages, in megabytes.

Memory
- **Memory Used %**- Traces the percent of memory used.

Subscription Mem Used % - Traces the percent of memory used by subscriptions.

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

Time Settings
By default, the time range end point is the current time. To change the time range, click the **Time Settings** and either:
- choose a **Time range** from 5 Minutes to 7 Days in the drop-down menu.
- specify begin/end dates using the calendar .
- specify begin/end time using the clock .

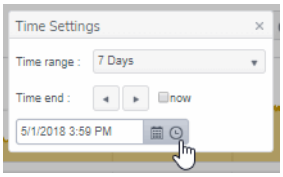

Toggle forward/backward in the trend graph per the period you choose (from the **Time range** drop-down menu) using arrows .

Restore settings to current time by selecting **now** .

Routers - Sensors

This tabular display contains environmental sensor metrics for a selected message router. Use this display to find out the type, name, value, and status of the sensors.

Select a router from the drop-down menu. Note that display does not show data for VMRs as it only applies to message routers.

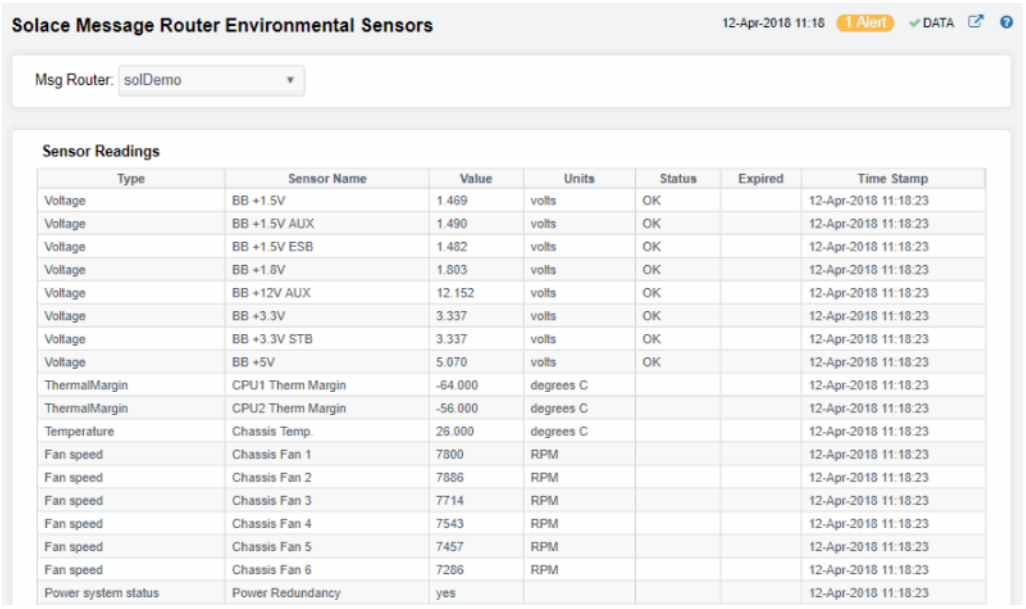

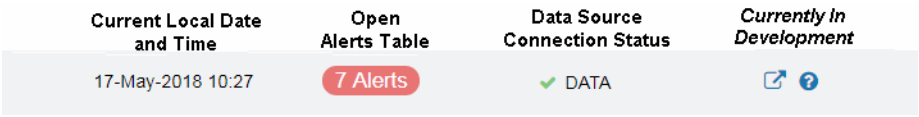

Sensor Readings
Each row in the table is a different sensor on the message router.

Type See vendor documentation for details.

Sensor Name The name of the sensor.

Value Lists the value of the sensor.

Units Lists the unit of measure for the sensor.

Status The current status of the sensor.

Expired When checked, performance data about the sensor has not been received within the time specified (in seconds) in the **$solRowExpirationTime** field in the **conf\rtvapm_solmon.properties** file. The **$solRowExpirationTimeForDelete** field allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response from the sensor. To view/edit the current values, modify the following lines in the **.properties** file:

```
# Metrics data are considered expired after this number of seconds
#
collector.sl.rtview.sub=$solRowExpirationTime:45
collector.sl.rtview.sub=$solRowExpirationTimeForDelete:3600
```

In the example above, the **Expired** check box would be checked after 45 seconds, and the row would be removed from the table after 3600 seconds.

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

Routers - Provisioning

This display shows provisioning metrics for a single message router. Use this to see the host, platform, chassis, memory, redundancy and fabric data for a specific message router.

Select a router from the drop-down menus.

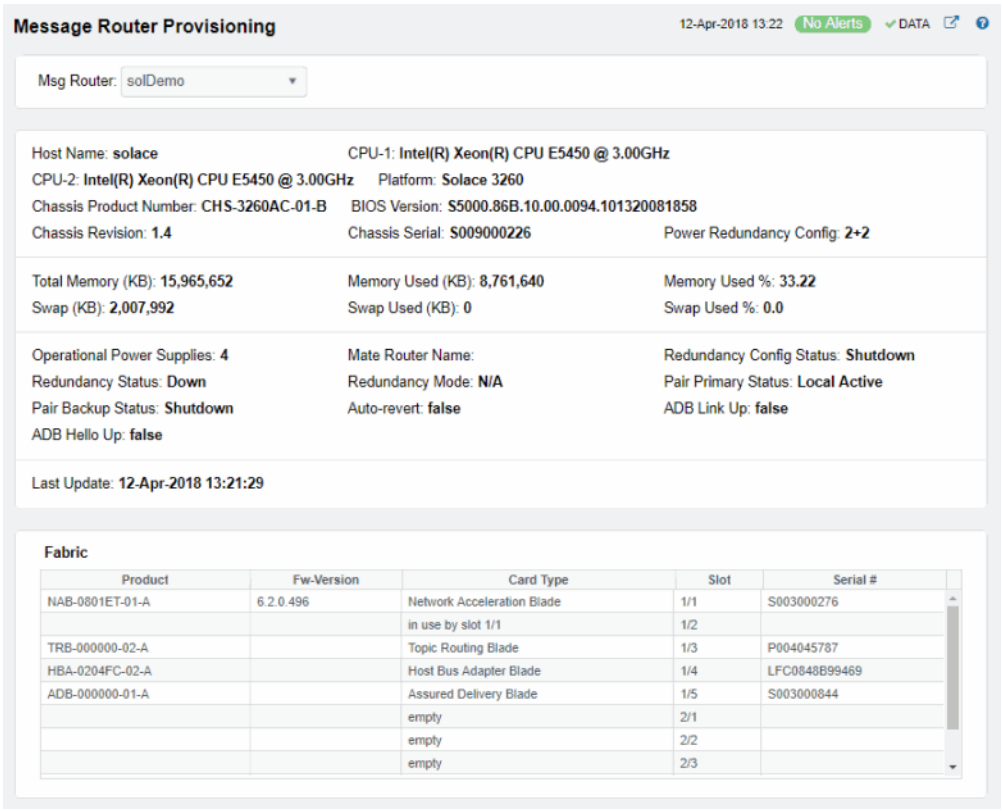

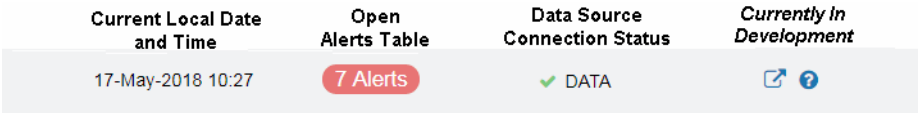

Host Name	The name of the host.
Platform	The platform on which the message router is running.
Chassis Product #	The product number of the chassis in which the router is contained.
Chassis Revision #	The revision number of the chassis.
Chassis Serial #	The serial number of the chassis.
Power Configuration	The power configuration used by the chassis.
Operational Power Supplies	The number of available power supplies that are operational on the chassis.
CPU 1	The name of the central processing unit (CPU 1) used by the message router.

CPU 2	The name of the central processing unit (CPU 2) used by the message router.	
BIOS	The basic input/output system used by the chassis.	
Memory (KB)		
	Physical	Lists the **Total** amount, the **Free** amount, the **Used** amount, and the **Used %** of physical memory.
	Swap	Lists the **Total** amount, the **Free** amount, the **Used** amount, and the **Used %** of swap memory.

Redundancy
These fields describe a fault tolerant pair of message routers.

	Mate Router Name	If redundancy is configured, this field lists the redundant router name (mate router name).
	Configuration Status	The status of the configuration for the backup message router.
	Redundancy Status	The status of the redundant message router.
	Redundancy Mode	Refer to Solace documentation for more information.
	Primary Status	The status of the primary message router.
	Backup Status	Refer to Solace documentation for more information.
	Auto-Revert	Refer to Solace documentation for more information.
	ADB Link Up	This check box is checked if a message router is set up to use guaranteed messaging and an Assured Delivery Blade (ADB) is set up and working correctly.
	ADB Hello Up	Refer to Solace documentation for more information.
Fabric		
	Slot	Displays the slot number on the network switch.
	Card Type	The type of card connected to the particular slot.
	Product	The product associated with the particular slot.
	Serial #	The serial number of the product.
	Fw-Version	The firmware version of the product.

Routers - Interface

This display lists all network interfaces on a selected message router, the status and in/out throughput per second for each network interface, as well as detailed metrics for a selected network interface.

Select a router and interface from the drop-down menus. Each row in the table is a different network interface. Double-click one to trace its current and historical performance data in the trend graph (bytes in/out and packets in/out per second).

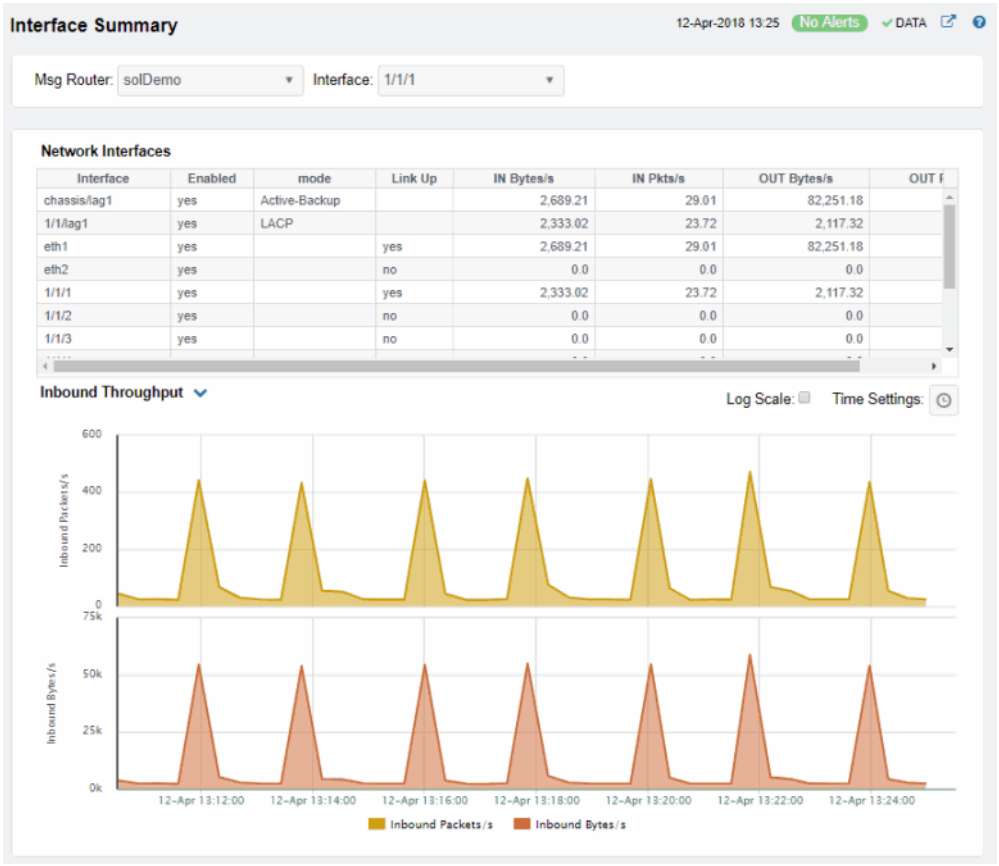

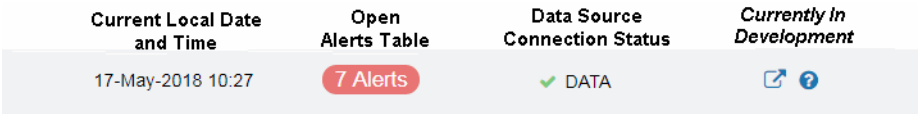

Interface	The name of the network interface.
Enabled	Displays whether or not the network interface is enabled.
mode	Describes how the interface is configured to support networking operations.
Link Up	Indicates whether the interface is electrically signaling on the transmission medium.
IN Bytes/sec	The number of bytes per second contained in incoming messages.
IN Pkts/sec	The number of incoming packets per second.
OUT Bytes/ sec	The number of bytes per second contained in the outgoing messages.
OUT Pkts/sec	The number of outgoing packets per second.

Trend Graphs

Inbound Pkts/ sec	Traces the number of incoming packets per second.
Outbound Bytes/sec	Traces the number of bytes per second contained in the incoming messages.
Log Scale	Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
Time Settings	By default, the time range end point is the current time. To change the time range, click the **Time Settings** and either: • choose a **Time range** from 5 Minutes to 7 Days in the drop-down menu. • specify begin/end dates using the calendar . • specify begin/end time using the clock . 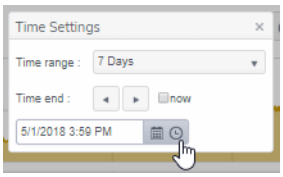 Toggle forward/backward in the trend graph per the period you choose (from the **Time range** drop-down menu) using arrows . Restore settings to current time by selecting **now** .

Routers - Message Spool

This display shows operational status and message spool metrics (if spooling is enabled on the message router) for one or all message routers. Refer to Solace documentation for details about data in this display.

Select a router from the drop-down menu.

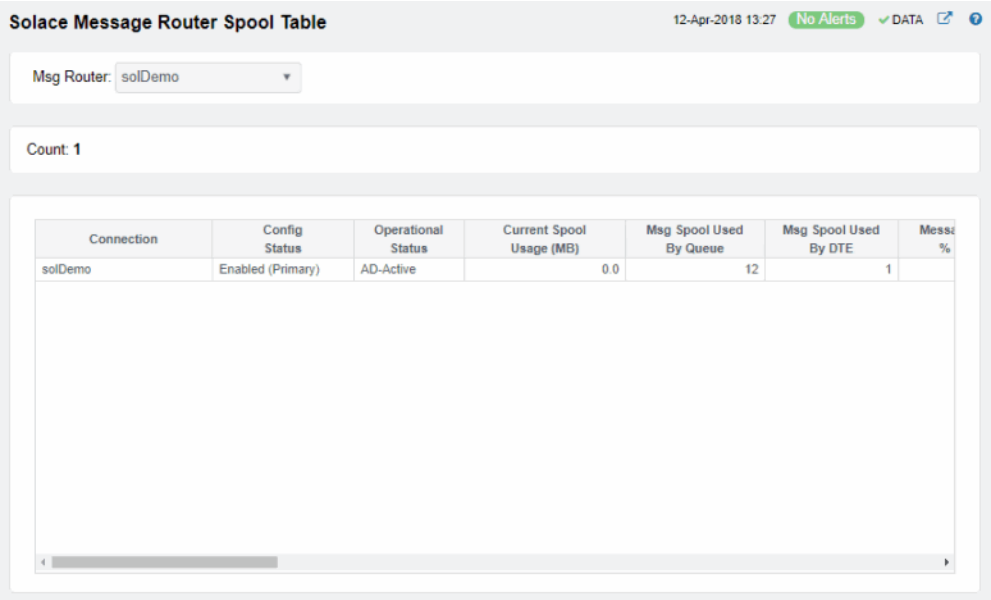

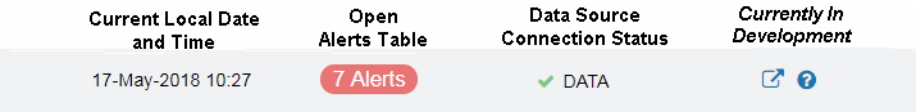

Count	The number of message routers that are using spooling in the table.
Connection	The name of the message router.
Config Status	The status of the connection's configuration.
Operational Status	The operational status of the spool on the message router.
Current Spool Usage (MB)	The current amount of spool used in megabytes on the message router (calculated by summing spool used for each endpoint).
Msg Spool Used By Queue	The amount of spool used by the queue.
Msg Spool Used By DTE	The amount of spool used by DTE.
Message Count % Utilization	The percentage of total messages that use the message spool.
Delivered UnAcked Msgs % Utilization	The percentage of messages delivered via the spool that have not been acknowledged.
Ingress Flow Count	The current incoming flow count.
Ingress Flows Allowed	The total number of incoming flows allowed.
Queue/Topic Subscriptions Used	The number of queue/topic subscriptions used.
Max Queue/Topic Subscriptions	The maximum number of queue/topic subscriptions available.

Sequenced Topics Used	The number of sequenced topics used.
Max Sequenced Topics	The maximum number of sequenced topics available.
Spool Files Used	The number of spool files used.
Spool Files Available	The maximum number of spool files available.
Spool Files % Utilization	The percentage of available spool files that have been used.
Active Disk Partition % Usage	The percentage of available active disk partition that has been used.
Standby Disk Partition % Usage	The percentage of available standby disk partition that has been used.
Disk Usage Current (MB)	The current amount of spool disk usage in megabytes.
Disk Usage Max (MB)	The maximum amount of available spool disk usage in megabytes.
Transacted Sessions Used	The current number of transacted sessions.
Transacted Sessions Max	The maximum number of transacted sessions allowed.
Transacted Session Count % Utilization	The percentage of allowable transacted sessions that have been used.
Transacted Session Resource % Utilization	The percentage of allowable transacted session resources that have been used.

Expired	When checked, performance data about the message router has not been received within the time specified (in seconds) in the **$solRowExpirationTime** field in the **conf\rtvapm_solmon.properties** file. The **$solRowExpirationTimeForDelete** field allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response from the message router. To view/edit the current values, modify the following lines in the **.properties** file:
	`# Metrics data are considered expired after this number of seconds`
	`#`
	`collector.sl.rtview.sub=$solRowExpirationTime:45`
	`collector.sl.rtview.sub=$solRowExpirationTimeForDelete:3600`
	
	In the example above, the **Expired** check box would be checked after 45 seconds, and the row would be removed from the table after 3600 seconds.

Neighbors

These displays provide detailed data and statuses for CSPF neighbor message routers. Check trends on network traffic among CSPF neighbors. Displays in this View are:

- "Neighbors - Table": View metrics for Solace neighbor message routers that use the Content Shortest Path First (CSPF) routing protocol to determine the shortest path in which to send messages from one message router to another message router in the Solace network.
- "Neighbors - Diagram": Topological view of CSPF Neighbors that shows message router connections and status of servers (Active/Inactive).
- "Neighbors - Summary": View detailed performance metrics for a single Solace neighbor message router that uses the CSPF routing protocol.

Neighbors - Table

This tabular display shows Content Shortest Path First (CSPF) "neighbor" metrics for a message router. Select a router from the drop-down menu. View metrics for a Solace neighbor message router that uses the CSPF routing protocol to determine the least cost path in which to send messages from one message router to another message router in the Solace network.

By default, a subset of available metrics is shown. Use **More Columns**/**Less Columns** to toggle to the complete set of metrics available (and back to the subset).

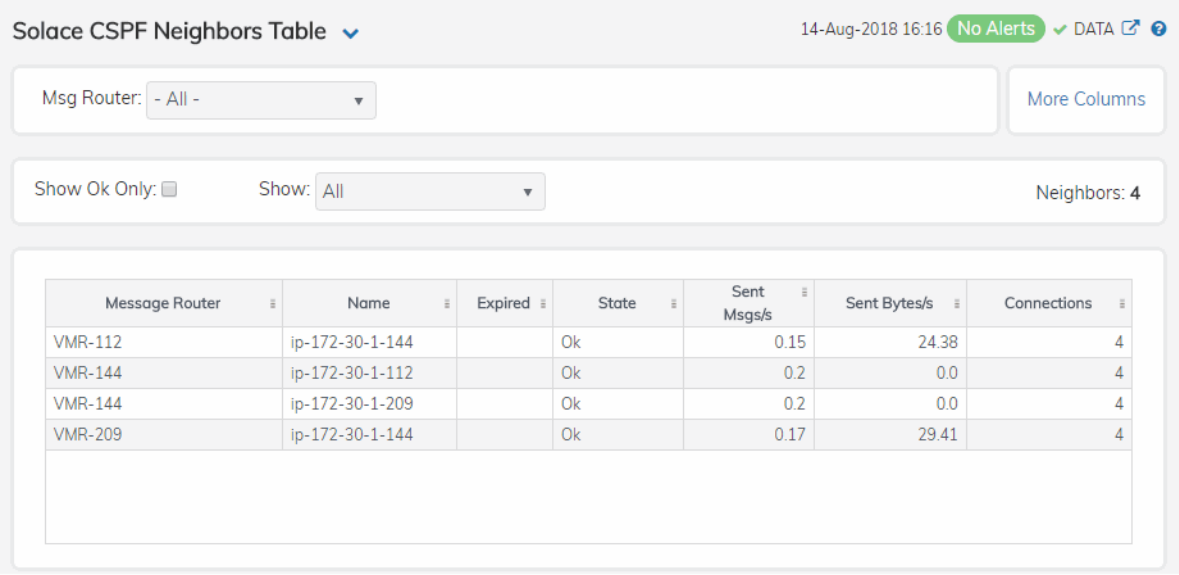

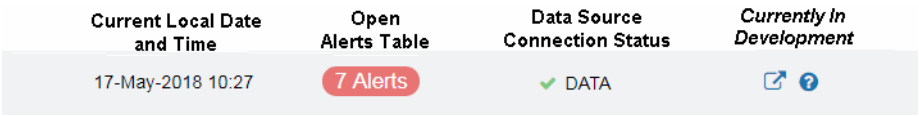

Neighbor Count:	The number of neighbor message routers connected to the selected **Msg Router**.	
Show:	**OK**	Select to *only* show neighbor message routers that are connected (**State** is **OK**). By default, this option is not selected (all neighbor message routers are shown.
	Expired	Select to show *both* expired and non-expired neighbor message routers. By default, this option is not selected (only non-expired neighbor message routers are shown).

Table:
Each table row is a different neighbor message router.

Message Router	The name of the neighbor message router.
State	The current state of the message router.
Up Time	The amount of time the message router has been up and running.
Connections	The number of connections.
Link Cost Actual	Refer to Solace documentation for more information.
Link Cost Configured	Refer to Solace documentation for more information.
Data Port	Refer to Solace documentation for more information.

Expired	When checked, performance data about the message router has not been received within the time specified (in seconds) in the **$solRowExpirationTime** field in the **conf\rtvapm_solmon.properties** file. The **$solRowExpirationTimeForDelete** field allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response from the message router. To view/edit the current values, modify the following lines in the **.properties** file:
	`# Metrics data are considered expired after this number of seconds`
	`#`
	`collector.sl.rtview.sub=$solRowExpirationTime:45`
	`collector.sl.rtview.sub=$solRowExpirationTimeForDelete:3600`
	
	In the example above, the **Expired** check box would be checked after 45 seconds, and the row would be removed from the table after 3600 seconds.

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

Neighbors - Diagram

Use this topology display to monitor the health of network components: Solace message routers, VMRs and servers. Quickly identify router neighbors, servers that are inactive and which resources their performance impacts. Drag and drop objects to arrange them on the screen (doing so does not logically impact the Solace message routers, VMRs and servers).

Each object is a Solace message router, VMR or server. Each are labeled with their name and color coded as follows:

- Red indicates that the object has one or more alerts in a critical state.
- Yellow indicates that the object has one or more alerts in a warning state.
- Green indicates that there are no alerts on the object.

Gray indicates that the object is off-line.

Mouse-over objects to see their host IP address.

Right-click on VMR objects and select **Open VMR UI** to open the Solace VMR login web page.

Save: Saves the arrangement of the objects.

Restore: Returns objects to their previous positions.

Layout: Toggles between two types of layouts. One layout positions objects to the right so you might scroll in that direction to see them. The other layout pulls the objects close together to the left, vertically and in hierarchical order.

Look at the miniature view in (upper left) to see all objects in either layout. Zoom in on an area in the topology by clicking it in the miniature view.

To monitor network bridges and VPNs, see the "Bridges - Diagram".

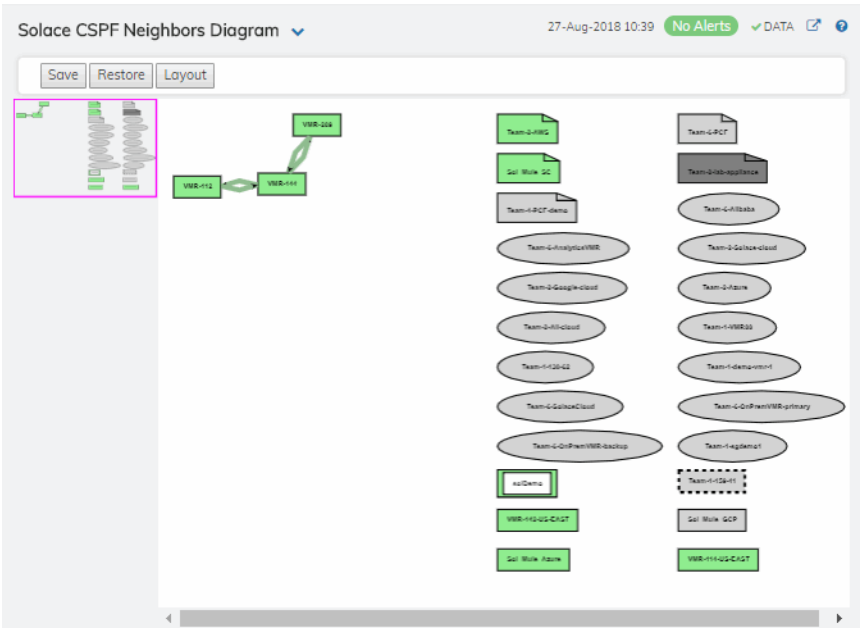

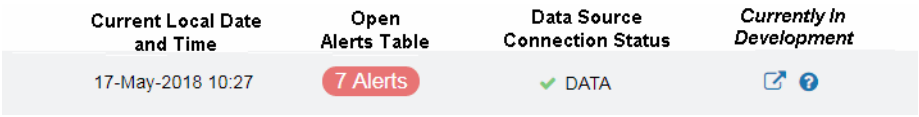

Neighbors - Summary

View neighbor message router current configuration details and message throughput rates.

Select a message router and a neighbor message router from the drop down menus. Check message throughput rates to the neighbor message router, as well as neighbor **Up Time**, **State**, **Data Port**, number of connections and link costs.

The trend graph traces the current and historical message throughput (**Data**, **Control**, **Discards** and **Total**).

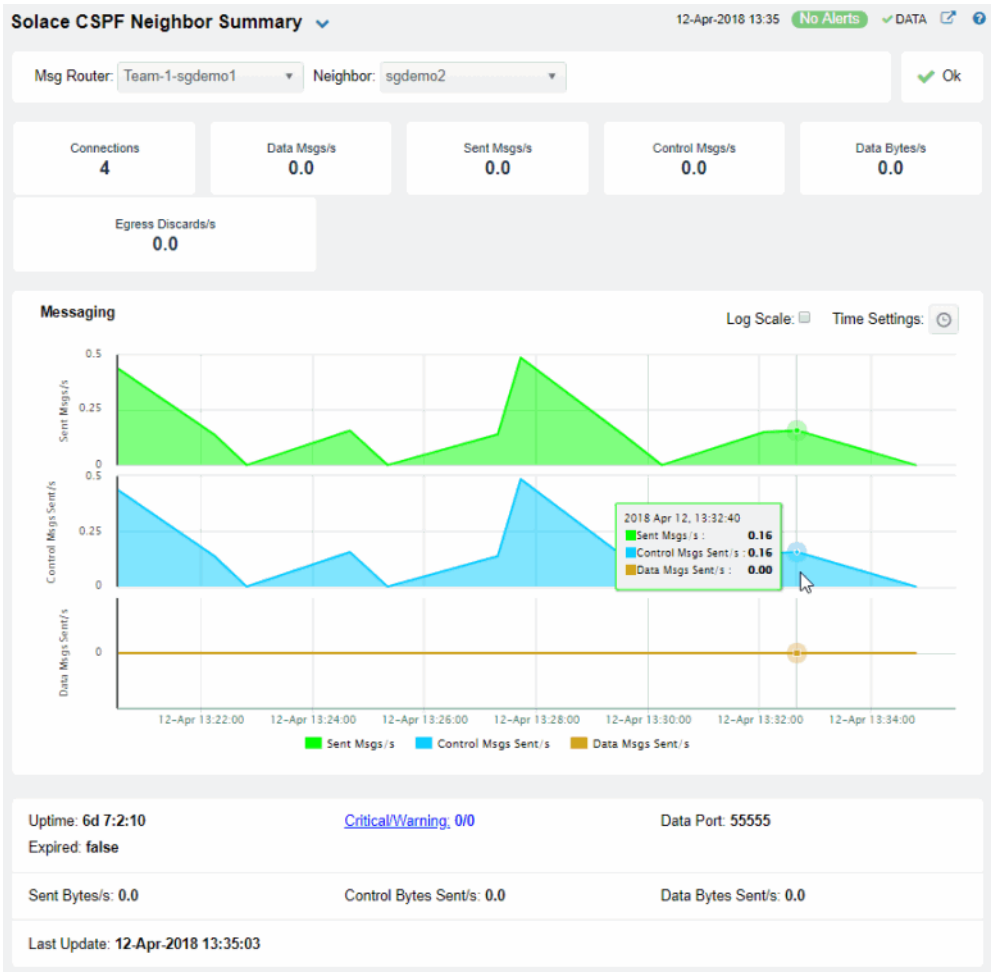

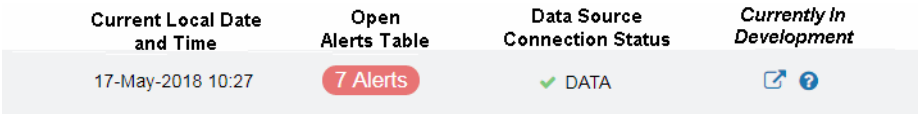

Neighbor:	Select the neighbor message router for which you want to show data in the display.
Connections	The current number of connections.
Data Msgs/s	Refer to Solace documentation for more information.
Sent Msgs/s	Refer to Solace documentation for more information.
Control Msgs/ s	Refer to Solace documentation for more information.
Data Bytes/s	Refer to Solace documentation for more information.

Egress Discards/s	The total number of discarded messages sent from the selected **Msg Router** to the selected **Neighbor** message router since the message router was last started.

Trend Graphs
Traces the rates of messages sent from the selected **Msg Router** to the selected **Neighbor** message router.

Sent Msgs/s	Refer to Solace documentation for more information.
Control Msgs/s	Refer to Solace documentation for more information.
Discards/s	Traces the number of discarded messages sent, per second, from the selected **Msg Router** to the selected **Neighbor** message router.
Log Scale	Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
Time Settings	By default, the time range end point is the current time. To change the time range, click the **Time Settings** and either: • choose a **Time range** from 5 Minutes to 7 Days in the drop-down menu. • specify begin/end dates using the calendar . • specify begin/end time using the clock . 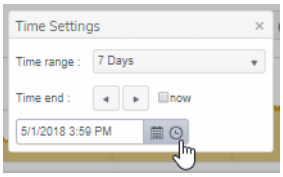 Toggle forward/backward in the trend graph per the period you choose (from the **Time range** drop-down menu) using arrows . Restore settings to current time by selecting **now** .

VPNs

You can view data for all VPNs configured on a specific message router in heatmap, table, or grid formats, or you can view data for a single VPN. Displays in this View are:

- "VPNs - Heatmap" on page 84: A color-coded heatmap view of the current status of all VPNs configured on a specific message router.
- "VPNs - Table" on page 88: A tabular view of all available data for all VPNs configured on a specific router.
- "VPNs - Summary" on page 92: Current and historical metrics for a single VPN.

VPNs - Heatmap

View the status of all VPNs configured on a specific message router in a heatmap format, which allows you to quickly identify VPNs with critical alerts. Each rectangle in the heatmap represents a VPN. The rectangle color indicates the alert state for each VPN.

Select a message router from the **Msg Router** drop-down menu and select a metric from the **Metric** drop-down menu. Use the **Operational** check-box ☑ to include or exclude non-operational VPNs in the heatmap. By default, this display shows **Alert Severity**, but you can mouse over a rectangle to see additional metrics. Drill-down and investigate by clicking a rectangle in the heatmap to view details for the selected application in the "VPNs - Summary" display.

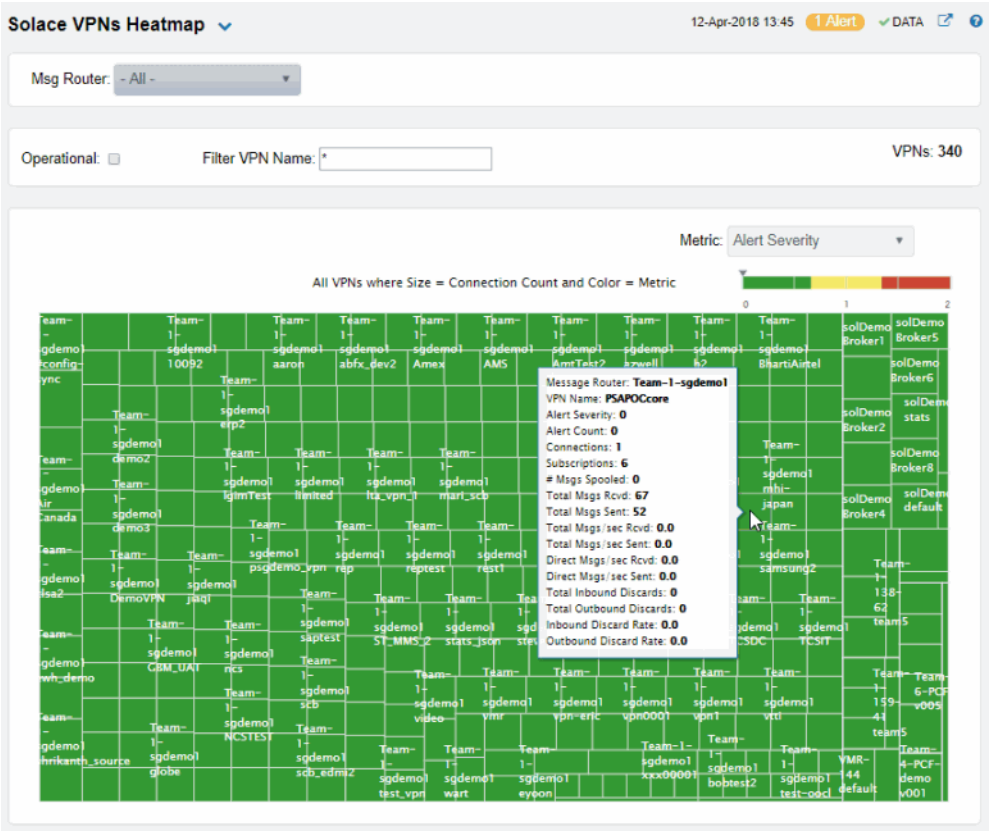

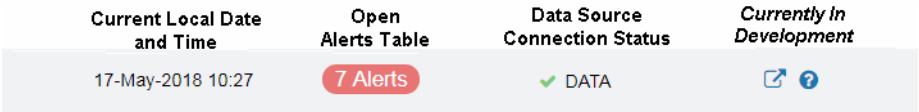

Operational When checked, only shows operational message routers.

Filter VPN Name Enter a string to show only VPNs with this string in their name.

Metric Choose a metric to view in the display.

Alert Severity

Visually displays the level at which the VPN has or has not exceeded its alarm level threshold. Values range from **0** - **2**, as indicated in the color gradient bar, where **2** is the highest Alert Severity:

Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.

Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.

Green indicates that no metrics have exceeded their alert thresholds.

Alert Count

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

Connections

The total number of connections. The color gradient 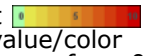 bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **SolVpnConnectionCountHigh**. The middle value in the gradient bar indicates the middle value of the range.

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

Subscriptions

The total number of subscriptions. The color gradient 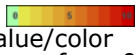 bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **SolVpnSubscriptionCountHigh**. The middle value in the gradient bar indicates the middle value of the range.

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

 # Msgs Spooled

The total number of spooled messages. The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **SolMsgRouterPendingMsgsHigh**. The middle value in the gradient bar indicates the middle value of the range.

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

Total Msgs Rcvd

The total number of received messages. The color gradient 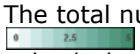 bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of messages received in the heatmap. The middle value in the gradient bar indicates the average count.

The **Auto** flag does not impact this metric.

Total Msgs Sent

The total number of sent messages. The color gradient 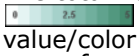 bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of messages sent in the heatmap. The middle value in the gradient bar indicates the average count.

The **Auto** flag does not impact this metric.

Total Msgs/ sec Rcvd

The number of messages received per second. The color gradient 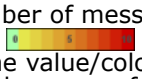 bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **SolVpnInboundMsgRateHigh**. The middle value in the gradient bar indicates the middle value of the range.

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

Total Msgs/ sec Sent

The number of messages sent per second. The color gradient 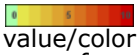 bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **SolVpnOutboundMsgRateHigh**. The middle value in the gradient bar indicates the middle value of the range.

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

Total Bytes/ sec Rcvd

The number of bytes contained in messages received per second. The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **SolVpnInboundByteRateHigh**. The middle value in the gradient bar indicates the middle value of the range.

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

Total Bytes/ sec Sent

The number of bytes contained in direct messages sent per second. The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **SolMsgRouterOutboundByteRateHigh**. The middle value in the gradient bar indicates the middle value of the range.

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

Direct Msgs/ sec Rcvd

The number of direct messages received per second. The color gradient 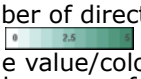 bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the average number of direct messages received per second in the heatmap. The middle value in the gradient bar indicates the average count.

The **Auto** flag does not impact this metric.

Direct Msgs/ sec Sent

The number of direct messages sent per second in the heatmap rectangle. The color gradient 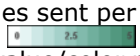 bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the average number of direct messages sent per second in the heatmap. The middle value in the gradient bar indicates the average count.

The **Auto** flag does not impact this metric.

Total Inbound Discards

The total number of discarded inbound messages in the heatmap rectangle. The color gradient 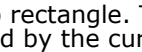 bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of discarded inbound messages in the heatmap. The middle value in the gradient bar indicates the average count.

The **Auto** flag does not impact this metric.

Total Outbound Discards	The total number of discarded outbound messages in the heatmap rectangle. The color gradient 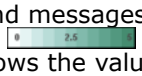 bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of discarded outbound messages in the heatmap. The middle value in the gradient bar indicates the average count. The **Auto** flag does not impact this metric.
Inbound Discard Rate	The number of discarded inbound messages per second in the heatmap rectangle. The color gradient 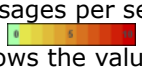 bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **SolVpnInboundDiscardRateHigh**. The middle value in the gradient bar indicates the middle value of the range. When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.
Outbound Discard Rate	The number of discarded outbound messages per second in the heatmap rectangle. The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **SolVpnOutboundDiscardRateHigh**. The middle value in the gradient bar indicates the middle value of the range. When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

VPNs - Table

View data shown in the "VPNs - Heatmap" display, as well as additional details, in a tabular format. Use this display to view all available data for each VPN associated with a specific message router.

By default, a subset of available metrics is shown. Use **More Columns**/**Less Columns** to toggle to the complete set of metrics available (and back to the subset).

Select a message router from the **Msg Router** drop-down menu. Each table row is a different VPN associated with the router. Click a column header to sort column data in numerical or alphabetical order.

Sort data in numerical or alphabetical order on column headers. Use the check-box ☑ to include / exclude non-operational VPNs. Use the **Show** drop-down to see **All** VPNs, **Expired Only** or **Unexpired Only**. Enter a string to show only VPNs with this string in their name.

Double-click a row to drill-down and investigate in the "VPNs - Summary" display.

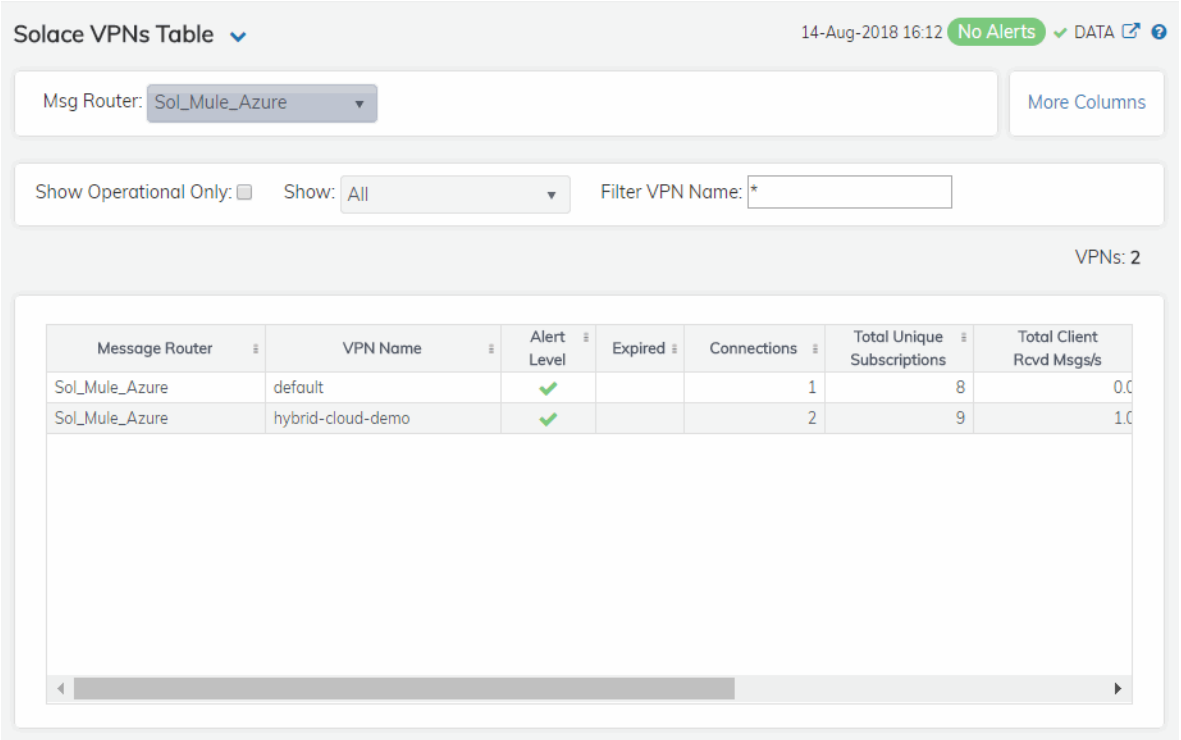

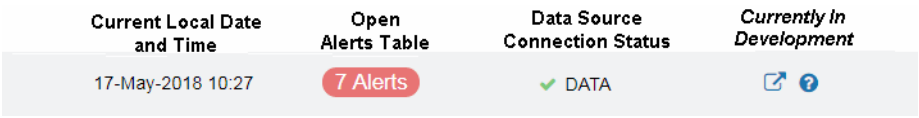

Message Router	The name of the message router.
VPN Name	The name of the VPN.
Alert Level	The maximum level of alerts in the row: Red indicates that one or more metrics exceeded their ALARM LEVEL threshold. Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold. Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	The total number of active alerts for the VPN.
Connections	The total number of connections for the VPN.
Operational	When checked, this status indicates that the VPN is enabled and is operating normally.
Total Unique Subscriptions	The total number of unique subscriptions to the VPN.
Total Client Messages Rcvd	The total number of messages received from clients connected to the VPN.
Total Client Messages Sent	The total number of messages sent to clients connected to the VPN.

Total Client Bytes Rcvd	The total number of bytes contained in messages received from clients connected to the VPN.
Total Client Bytes Sent	The total number of bytes contained in messages sent to clients connected to the VPN.
Total Client Msgs/sec Rcvd	The total number of messages received per second from clients connected to the VPN.
Total Client Msgs /sec Sent	The total number of messages sent per second to clients connected to the VPN.
Total Client Bytes/sec Rcvd	The total number of bytes contained in messages received per second from clients connected to the VPN.
Total Client Bytes/sec Sent	The total number of bytes contained in messages sent per second to clients connected to the VPN.
Client Direct Msgs Rcvd	The total number of direct messages received from clients connected to the VPN.
Client Direct Msgs Sent	The total number of direct messages sent to clients connected to the VPN.
Client Direct Bytes Rcvd	The total number of bytes contained in direct messages received from clients connected to the VPN.
Client Direct Bytes Sent	The total number of bytes contained in direct messages sent to clients connected to the VPN.
Client Direct Msgs/sec Rcvd	The total number of direct messages received per second from clients connected to the VPN.
Client Direct Msgs/sec Sent	The total number of direct messages sent per second to clients connected to the VPN.
Client Direct Bytes/sec Rcvd	The total number of bytes contained in the direct messages received per second from clients connected to the VPN.
Client Direct Bytes/sec Sent	The total number of bytes contained in the direct messages sent per second to clients connected to the VPN.
Client NonPersistent Msgs Rcvd	The total number of non-persistent messages received from clients connected to the VPN.
Client NonPersistent Msgs Sent	The total number of non-persistent messages sent to clients connected to the VPN.
Client NonPersistent Bytes Rcvd	The total number of bytes contained in the non-persistent messages received from clients connected to the VPN.
Client NonPersistent Bytes Sent	The total number of bytes contained in the non-persistent messages sent per second to clients connected to the VPN.
Client NonPersistant Msgs/ sec Rcvd	The total number of non-persistent messages received per second from clients connected to the VPN.
Client NonPersistent Msgs/ sec Sent	The total number of non-persistent messages sent per second to clients connected to the VPN.
Client NonPersistant Bytes/ sec Rcvd	The total number of bytes contained in the non-persistent messages received per second from clients connected to the VPN.
Client NonPersistent Bytes/ sec Sent	The total number of bytes contained in the non-persistent messages sent per second to clients connected to the VPN.
Client Persistent Msgs Rcvd	The total number of persistent messages received from clients connected to the VPN.
Client Persistent Msgs Sent	The total number of persistent messages sent to clients connected to the VPN.

Client Persistent Bytes Rcvd	The total number of bytes contained in persistent messages received from clients connected to the VPN.
Client Persistent Bytes Sent	The total number of bytes contained in persistent messages sent to clients connected to the VPN.
Client Persistent Msgs/sec Rcvd	The total number of persistent messages received per second from clients connected to the VPN.
Client Persistent Msgs/sec Sent	The total number of persistent messages sent per second to clients connected to the VPN.
Client Persistent Bytes/sec Rcvd	The total number of bytes contained in the persistent messages received per second from clients connected to the VPN.
Client Persistent Bytes/sec Sent	The total number of bytes contained in the persistent messages sent per second to clients connected to the VPN.
Total In Discards	The total number of discarded incoming messages.
Total In Discards/sec	The number of discarded incoming messages per second.
Total Out Discards	The total number of discarded outgoing messages.
Total Out Discards/sec	The number of discarded outgoing messages per second.
Max Spool Usage (MB)	The maximum amount of disk storage (in megabytes) that can be consumed by all spooled message on the VPN.
Authentication Type	The defined authentication type on the VPN.

Expired	When checked, performance data about the VPN has not been received within the time specified (in seconds) in the **$solRowExpirationTime** field in the **conf\rtvapm_solmon.properties** file. The **$solRowExpirationTimeForDelete** field allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response from the VPN. To view/edit the current values, modify the following lines in the **.properties** file:
	`# Metrics data are considered expired after this number of seconds`
	`#`
	`collector.sl.rtview.sub=$solRowExpirationTime:45`
	`collector.sl.rtview.sub=$solRowExpirationTimeForDelete:3600`
	
	In the example above, the **Expired** check box would be checked after 45 seconds, and the row would be removed from the table after 3600 seconds.

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

VPNs - Summary

Select a message router and a VPN to view details about alerts, connections/destinations, incoming messages and outgoing/pending messages for the VPN.

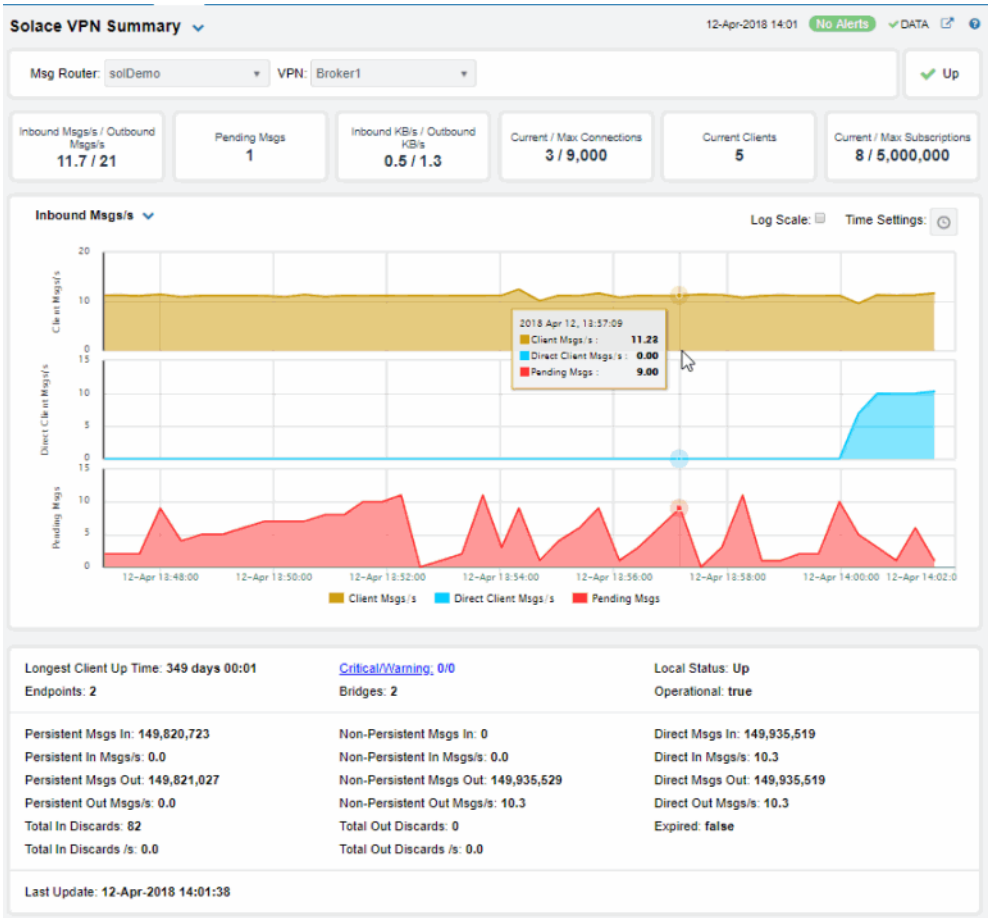

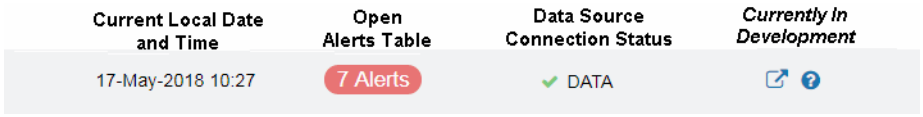

Alerts	Red indicates that one or more metrics exceeded their ALARM LEVEL threshold. Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold. Green indicates that no metrics have exceeded their alert thresholds.
Up	
Inbound/Outbound Msgs/s	The number of inbound/outbound messages per second.
Pending Msgs	The number of pending messages.
Inbound/Outbound KB/s	The number of inbound/outbound messages in KBs per second.

Current/Max Connections	The total number of current connections / maximum number of supported connections for the VPN.
Current Clients	The number of connected clients.
Current/Max Subscriptions	The total number of current subscribers and maximum number of supported subscribers for the VPN.

Inbound Msgs/s Trend Graphs
Traces the sum of inbound message processing for the selected VPN.
- **Pending Msgs**: The number of pending messages for the VPN.
- **Client Msgs/sec**: The rate of incoming messages (per second) from client.
- **Direct Client Msgs/sec**: The rate of direct incoming messages (per second) from the direct client.

Log Scale	Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
Time Settings	By default, the time range end point is the current time. To change the time range, click the **Time Settings** and either: • choose a **Time range** from 5 Minutes to 7 Days in the drop-down menu. • specify begin/end dates using the calendar . • specify begin/end time using the clock . 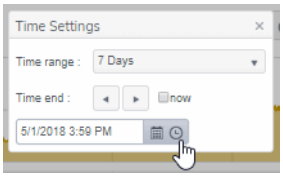 Toggle forward/backward in the trend graph per the period you choose (from the **Time range** drop-down menu) using arrows . Restore settings to current time by selecting **now** .
Longest Client Up Time	The number of days, hours and minutes for the longest, currently active, client connection.
Endpoints	The number of endpoints.
Persistent Msgs In	The total number of incoming persistent messages.
Persistent In Msgs/s	The number of incoming persistent messages per second.
Persistent Msgs Out	The total number of outgoing persistent messages.
Persistent Out Msgs/s	The number of outgoing persistent messages per second.
Total In Discards	The total number of incoming messages that were discarded.
Total In Discards/sec	The total number of incoming messages that were discarded, per second.
Critical/Warning	The number of critical alerts / warning alerts which also opens the **Alerts Table**.

Bridges	The number of bridges.
Non-Persistent Msgs In	The total number of incoming non-persistent messages.
Non-Persistent In Msgs/s	The number of incoming non-persistent messages per second.
Non-Persistent Msgs Out	The total number of outgoing non-persistent messages.
Non-Persistent Out Msgs/s	The number of outgoing non-existent messages per second.
Total Out Discards	The total number of outgoing messages that were discarded.
Total Out Discards/sec	The total number of outgoing messages that were discarded, per second.
Direct Msgs In	The total number of incoming direct messages.
Direct In Msgs/s	The number of incoming direct messages per second.
Direct Msgs Out	The total number of outgoing direct messages.
Direct Out Msgs/s	The number of outgoing direct messages per second.

Expired	When **true**, performance data about the VPN has not been received within the time specified (in seconds) in the **$solRowExpirationTime** field in the **conf\rtvapm_solmon.properties** file. The **$solRowExpirationTimeForDelete** field allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response from the VPN. To view/edit the current values, modify the following lines in the **.properties** file:
	`# Metrics data are considered expired after this number of seconds`
	`#`
	`collector.sl.rtview.sub=$solRowExpirationTime:45`
	`collector.sl.rtview.sub=$solRowExpirationTimeForDelete:3600`
	
	In the example above, the **Expired** check box would be checked after 45 seconds, and the row would be removed from the table after 3600 seconds.

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

Clients

These displays allow you to view the current and historical metrics for clients configured on a VPN. Displays in this View are:

- "Clients - Table": A tabular view of data for all clients configured on a VPN.
- "Client - Summary": Current and historical metrics for a single client configured on a VPN.

Clients - Table

View data for all clients configured on a VPN. Select a router and VPN from the drop-down menus. Each table row is a different VPN client connection.

By default, a subset of available metrics is shown. Use **More Columns**/**Less Columns** to toggle to the complete set of metrics available (and back to the subset).

Use the drop-down menus to show **All**, **Expired** or **Unexpired** clients as well as **All**, **Internal** or **Primary** clients (processes that run on the message router under the Solace OS). Enter a string for **Filter Client Name** to show only clients with this string in their name.

This display is populated by two caches, SolClientsStats and SolClients. SolClientsStats provides most of the data. SolClients provides the static data. If the SolClients cache encounters an issue the static fields in this display are blank.

Double-click a row to drill-down and investigate in the "Client - Summary" display.

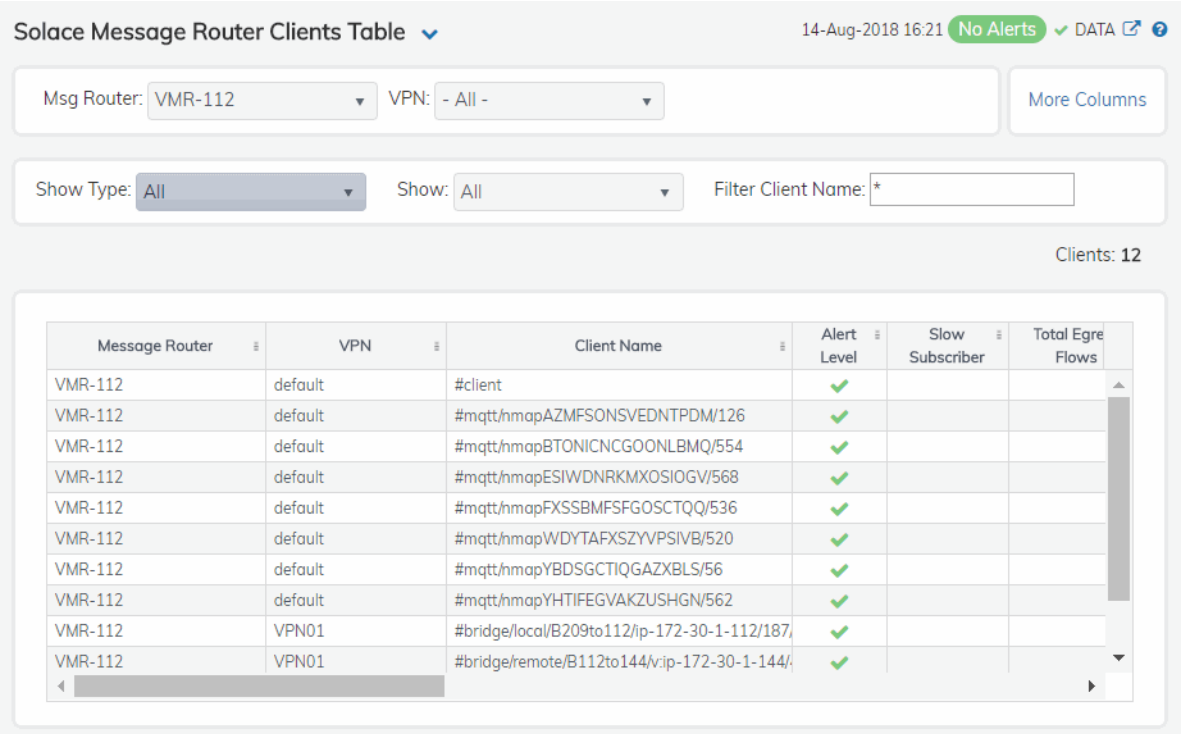

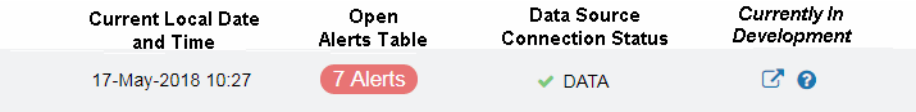

Message Router	Lists the name of the selected message router.
VPN	Lists the name of the selected VPN.
Client Name	The name of the client.
Alert Level	The maximum level of alerts in the row: Red indicates that one or more metrics exceeded their ALARM LEVEL threshold. Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold. Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	Total number of alerts for the client.

Slow Subscriber	This check box will be checked if the client consistently fails to consume their messages at the offered rate (which causes their egress queues to fill up).
Total Egress Flows	The total number of outgoing flows.
Total Ingress Flows	The total number of incoming flows.
Bind Requests	The number of bind requests made by the client.
Subscriptions	The total number of subscriptions.
Subscription Msgs Rcvd	The total number of messages received from subscriptions.
Subscription Msgs Sent	The total number of messages sent from subscriptions.
Type	Lists the type of alert.
Uptime	Lists the amount of time the client has been up and running.
Client ID	Lists the client ID.
Client UserName	Lists the user name for the client.
Client Address	The IP Address of the client.
Profile	The client profile that is assigned to the client.
ACL Profile	The access control list profile to which the client is assigned.
Description	Lists a description of the client.
Platform	Lists the platform of the client.
Software Version	The version of the platform.
Total Flows Out	The total number of outbound message flows for the client.
Total Flows In	The total number of inbound message flows for the client.
# Subscriptions	The number of subscribers connected to the client.
Add Sub Msgs Rcvd	The number of Add Subscription messages received.
Add Sub Msgs Sent	The number of Add Subscription Messages sent.
Already Exists Msgs Sent	Refer to Solace documentation for more information.
Assured Ctrl Msgs Rcvd	Refer to Solace documentation for more information.
Assured Ctrl Msgs Sent	Refer to Solace documentation for more information.
Total Client Msgs Rcvd	The total number of messages received by the client.
Total Client Msgs Sent	The total number of messages sent by the client.
Total Client Bytes Rcvd	The total number of bytes contained within the messages received by the client.
Total Client Bytes Sent	The total number of bytes contained within the messages sent by the client.
Total Client Msgs Rcvd/ sec	The total number of messages received per second by the client.

Total Client Msgs Sent/ sec	The total number of messages sent per second by the client.
Total Client Bytes Rcvd/ sec	The total number of bytes contained within the messages received per second by the client.
Total Client Bytes Sent/ sec	The total number of bytes contained within the messages sent per second by the client.
Ctl Bytes Rcvd	The number of control data bytes received by the client.
CTL Bytes Sent	The number of control data bytes sent by the client.
Ctl Msgs Rcvd	The number of control data messages received by the client.
Ctl Msgs Sent	The number of control data messages sent by the client.
Client Data Bytes Rcvd	The number of bytes contained within the data messages received by the client.
Client Data Bytes Sent	The number of bytes contained within the data messages sent by the client.
Client Data Msgs Rcvd	The number of data messages received by the client.
Client Data Msgs Sent	The number of data messages sent by the client.
Client Direct Msgs Rcvd	The number of direct messages received by the client.
Client Direct Msgs Sent	The number of direct messages sent by the client.
Client Direct Bytes Rcvd	The number of bytes contained within direct messages received by the client.
Client Direct Bytes Sent	The number of bytes contained within direct messages sent by the client.
Client Direct Msgs Rcvd/sec	The number of direct messages received per second by the client.
Client Direct Msgs Sent/ sec	The number of direct messages sent per second by the client.
Client Direct Bytes Rcvd/sec	The number of bytes contained within the messages received per second by the client.
Client Direct Bytes Sent/sec	The number of bytes contained within the messages sent per second by the client.
Client NonPersistent Msgs Rcvd	The number of non-persistent messages received by the client.
Client NonPersistent Msgs Sent	The number of non-persistent messages sent by the client.
Client NonPersistent Bytes Rcvd	The number of bytes contained within the non-persistent messages received by the client.
Client NonPersistent Bytes Sent	The number of bytes contained within the non-persistent messages sent by the client.
Client NonPersistent Msgs Rcvd/sec	The number of non-persistent messages received per second by the client.
Client NonPersistent Msgs Sent/sec	The number of non-persistent messages sent per second by the client.
Client NonPersistent Bytes Rcvd/sec	The number of bytes contained within the non-persistent messages received per second by the client

Client NonPersistent Bytes Sent/sec	The number of bytes contained within the non-persistent messages sent per second by the client
Client Persistent Msgs Rcvd	The number of persistent messages received by the client.
Client Persistent Msgs Sent	The number of persistent messages sent by the client.
Client Persistent Bytes Rcvd	The number of bytes contained within the persistent messages received by the client.
Client Persistent Bytes Sent	The number of bytes contained within the persistent messages sent by the client.
Client Persistent Msgs Rcvd/sec	The number of persistent messages received per second by the client.
Client Persistent Msgs Sent/sec	The number of persistent messages sent per second by the client.
Client Persistent Bytes Rcvd/sec	The number of bytes contained within the persistent messages received per second by the client.
Client Persistent Bytes Sent/sec	The number of bytes contained within the persistent messages sent per second by the client.
Denied Dup Clients	Refer to Solace documentation for more information.
Denied Subscribe Permission	The number of denied subscription requests due to improper permissions.
Denied Subscribe Topic-ACL	The number of denied subscriptions to topics due to the fact that the client requesting was not on the Access Control List.
Denied Unsubscribe Permission	The number of denied unsubscribe requests due to improper permissions.
Denied Unsubscribe Topic-ACL	The number of denied unsubscribe requests to topics due to the fact that the client requesting was not on the Access Control List.
DTO Msgs Rcvd	The number of Deliver-To-One messages received by the client.
Egress Compressed Bytes	The number of compressed bytes contained within outgoing messages.
Ingress Compressed Bytes	The number of compressed bytes contained within incoming messages.
Total Ingress Discards	The total number of discarded incoming messages.
Total Egress Discards	The total number of discarded outgoing messages.
Total Ingress Discards/ sec	The total number of discarded incoming messages per second.
Total Egress Discards/ sec	The total number of discarded outgoing messages per second.
Keepalive Msgs Rcvd	The number of Keepalive messages received by the client.
Keepalive Msgs Sent	The number of Keepalive messages sent by the client.
Large Msgs Rcvd	The number of large messages received by the client.
Login Msgs Rcvd	The number of login message received by the client.

Max Exceeded Msgs Sent	The number of responses sent by the client informing the connected message router(s) that the number of the message(s) sent exceeded the maximum allowed.
Not Enough Space Msgs Sent	The number of responses sent by the client informing the connected message router(s) that the size of the message(s) sent exceeded the maximum allowable size, or that the message caused the client's Local Spool Quota to exceed the maximum amount of space.
Not Found Msgs Sent	Refer to Solace documentation for more information.
Parse Error on Add Msgs Sent	Refer to Solace documentation for more information.
Parse Error on Remove Msgs Sent	Refer to Solace documentation for more information.
Remove Subscription Msgs Rcvd	The number of remove subscription requests received by the client.
Remove Subscription Msgs Sent	The number of remove subscription requests sent by the client.
Subscribe Client Not Found	The number of subscription requests for clients that were not found.
Unsubscribe Client Not Found	The number of unsubscribe requests for clients that were not found.
Update Msgs Rcvd	Refer to Solace documentation for more information.
Update Msgs Sent	Refer to Solace documentation for more information.

Expired	When checked, performance data about the client has not been received within the time specified (in seconds) in the **$solRowExpirationTime** field in the **conf\rtvapm_solmon.properties** file. The **$solRowExpirationTimeForDelete** field allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response from the client. To view/edit the current values, modify the following lines in the **.properties** file:
	`# Metrics data are considered expired after this number of seconds`
	`#`
	`collector.sl.rtview.sub=$solRowExpirationTime:45`
	`collector.sl.rtview.sub=$solRowExpirationTimeForDelete:3600`
	
	In the example above, the **Expired** check box would be checked after 45 seconds, and the row would be removed from the table after 3600 seconds.

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

Client - Summary

This display allows you to view the current and historical metrics for a single VPN client.

Select a router, VPN and client from the drop-down menus. You can view the **Client Type**, the **User Name**, the **Client ID**, the associated **Platform**, the current **Up Time**, and additional information specific to the client. You can also view the total number of incoming and outgoing messages, as well as the number of incoming and outgoing persistent, non-persistent, direct, and discarded messages.

This display is populated by two caches, SolClientsStats and SolClients. SolClientsStats provides most of the data. SolClients provides the static data. If the SolClients cache encounters an issue the graphic elements that have no data are replaced with **N/A**.

This display also includes a trend graph containing the current and historical incoming messages per second, outgoing messages per second, incoming direct messages per second, and outgoing direct messages per second.

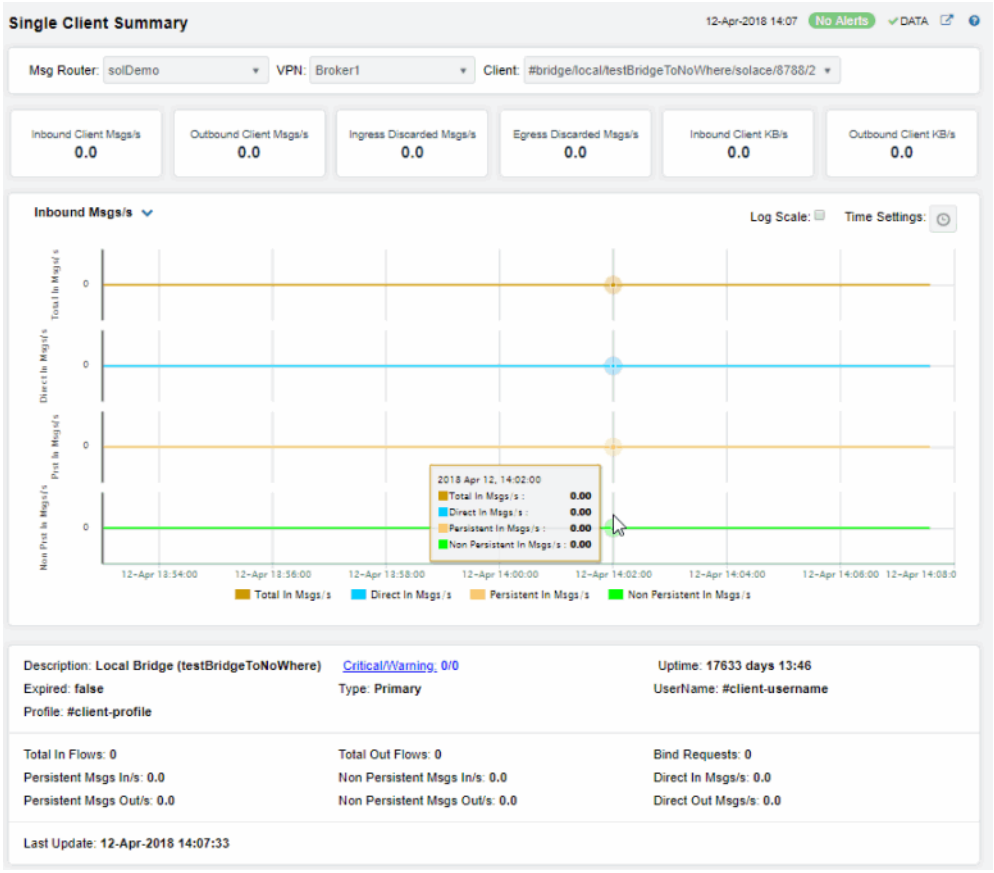

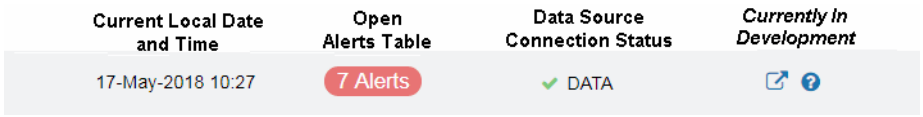

Inbound Client Msgs /sec	The number of incoming client messages per second.
Outbound Client Msgs /sec	The number of outgoing client messages per second.
Ingress Discarded Msgs / sec	The number of discarded ingress messages per second.
Egress Discarded Msgs / sec	The number of discarded egress messages per second.
Inbound Client KB/sec	The amount of incoming data from the client in KBs per second.
Outbound Client KB/sec	The amount of outgoing data for the client in KBs per second.

Trend Graphs
Traces the sum of message processing for the selected client.

- **Total In Msgs/sec**: The number of incoming messages (per second) for the client.
- **Dir-In Msgs/sec**: The number of incoming direct messages (per second) for the client.
- **Persistent In Msgs/sec**: The number of incoming persistent messages (per second) for the client.
- **Non Persitent In Msgs/sec**: The number of incoming non-persistent messages (per second) for the client.

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

Time Settings
By default, the time range end point is the current time. To change the time range, click the **Time Settings** and either:

- choose a **Time range** from 5 Minutes to 7 Days in the drop-down menu.
- specify begin/end dates using the calendar .
- specify begin/end time using the clock .

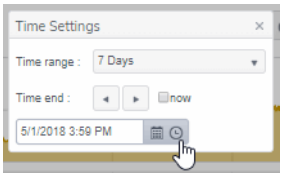

Toggle forward/backward in the trend graph per the period you choose (from the **Time range** drop-down menu) using arrows .
Restore settings to current time by selecting **now** .

Description
The description of the client.

Expired
When checked, performance data about the client has not been received within the time specified (in seconds) in the **$solRowExpirationTime** field in the **conf\rtvapm_solmon.properties** file. The **$solRowExpirationTimeForDelete** field allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response from the client. To view/edit the current values, modify the following lines in the **.properties** file:

```
# Metrics data are considered expired after this number of seconds
#
collector.sl.rtview.sub=$solRowExpirationTime:45
collector.sl.rtview.sub=$solRowExpirationTimeForDelete:3600
```

In the example above, the **Expired** check box would be checked after 45 seconds, and the row would be removed from the table after 3600 seconds.

Profile
The client's profile.

Total Ingress Flows
The number of inflows coming to the client.

Persistent Msgs In/sec	The number of persistent incoming messages per second.
Persistent Msgs Out/sec	The number of persistent outgoing messages per second.
Last Update	The date and time of the last data update.
Critical/Warning	The number of critical alerts / warning alerts which also opens the **Alerts Table**.
Non Persistent Msgs In/ sec	The number of non-persistent incoming messages per second.
NonPersistent Msgs Out/ sec	The number of non-persistent outgoing messages per second.
Uptime	If the VPN's **Local Status** is **Up**, this field displays the length of time that the VPN has been up and running.
Username	The client's user name.
Bind Requests	The number of bind requests received by the client.
Direct In Msgs /sec	The number of non-persistent incoming messages per second.
Direct Out Msgs /sec	The number of non-persistent outgoing messages per second.

Bridges

These displays provide process data for bridges configured on a VPN. Displays in this View are:

- "Bridges - Table": A tabular view of all available process performance data for all bridges configured on a VPN.
- "Bridges - Diagram": Topological view of Solace network bridges that shows bridge router connections and status.
- "Bridge - Summary": Current and historical metrics for a single bridge.

Bridges - Table

This display allows you to view data for all bridges configured for a VPN.

By default, a subset of available metrics is shown. Use **More Columns**/**Less Columns** to toggle to the complete set of metrics available (and back to the subset).

Select a router and VPN from the drop-down menus. Use the check-boxes ☑ to include / exclude **Enabled** and **Expired** bridges. Each table row is a different bridge.

Rows listing bridges that are disabled or expired display with a shaded background. Double-click a row to drill-down and investigate in the "Bridge - Summary" display.

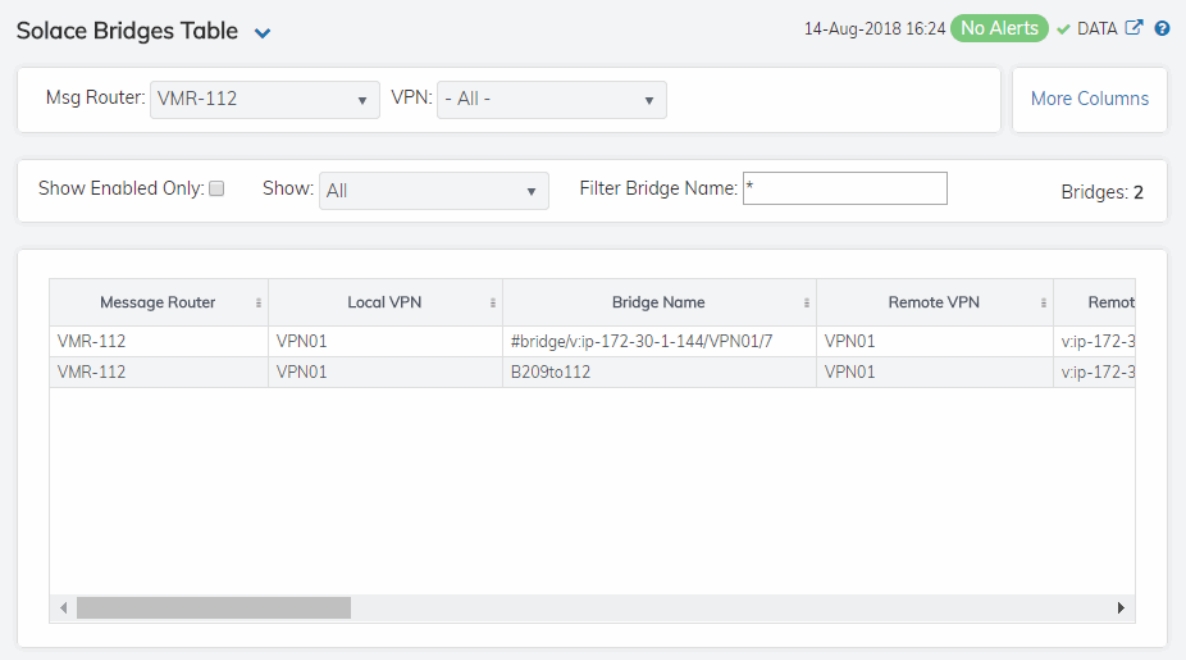

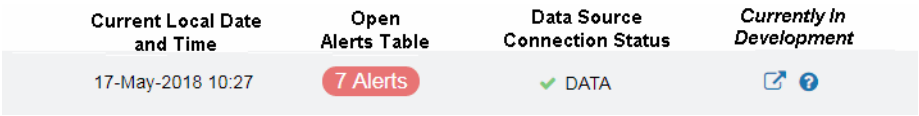

Message Router	Displays the name of the message router
Local VPN	The name of the local VPN.
Bridge Name	The name of the bridge.
Alert Level	The current level of alerts in the row. Red indicates that one or more metrics exceeded their ALARM LEVEL threshold. Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold. Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	The total number of active alerts for the process.
Remote VPN	The name of the remote VPN that is connected to the local VPN via the bridge.
Remote Router	The name of the remote router.
Admin State	Indicates whether the bridge has been administratively enabled (via SolAdmin or the command line interface).
Inbound Operational State	The current inbound operational status of the bridge. (The administrator can turn off a bridge's input or output for maintenance or other reasons.)

Outbound Operational State	The current outbound operational status of the bridge. (The administrator can turn off a bridge's input or output for maintenance or other reasons.)
Queue Operational State	The current operational status of the queue.
Connection Establisher	Indicates whether the administrator created and configured the bridge directly on the message router using SolAdmin or the command line interface, or indirectly from another message router.
Redundancy	Displays whether the bridge is the **primary** bridge, the **backup** bridge, the **static** bridge (default bridge used when no other bridge is available), or whether it is the only bridge available (**none**).
Uptime	The current amount of time in which the bridge has been up and running.
Client Name	The name of the client.
Connected Via Addr	The local IP address and port used for the bridge.
Connected Via Interface	The name of the network interface used for the bridge.
Client Direct Bytes Rcvd	The number of bytes contained within direct messages received by the client via the bridge.
Client Direct Bytes/sec Rcvd	The number of bytes contained within direct messages received per second by the client via the bridge.
Client Direct Bytes Sent	The number of bytes contained within direct messages sent by the client via the bridge.
Client Direct Bytes/sec Sent	The number of bytes contained within direct messages sent per second by the client via the bridge.
Client Direct Msgs/sec Rcvd	The number of bytes contained within direct messages received per second by the client via the bridge.
Client Direct Msgs Sent	The number of direct messages sent by the client via the bridge.
Client Direct Msgs/sec Sent	The number of direct messages sent per second by the client via the bridge.
Client NonPersistent Bytes Rcvd	The number of bytes contained within non-persistent messages received by the client via the bridge.
Client NonPersistent Bytes/sec Rcvd	The number of bytes contained within non-persistent messages received per second by the client via the bridge.
Client NonPersistent Bytes Sent	The number of bytes contained within non-persistent messages sent by the client via the bridge.
Client NonPersistent Bytes/sec Sent	The number of bytes contained within non-persistent messages sent per second by the client via the bridge.
Client NonPersistent Msgs Rcvd	The number of non-persistent messages received by the client via the bridge.
Client NonPersistent Msgs/sec Rcvd	The number of non-persistent messages received per second by the client via the bridge.
Client NonPersistent Msgs Sent	The number of non-persistent messages sent by the client via the bridge.
Client NonPersistent Msgs/sec Sent	The number of non-persistent messages sent per second by the client via the bridge.

Client Persistent Bytes Rcvd	The number of bytes contained within persistent messages received by the client via the bridge.
Client Persistent Bytes/ sec Rcvd	The number of bytes contained within persistent messages received per second by the client via the bridge.
Client Persistent Bytes Sent	The number of bytes contained within persistent messages sent by the client via the bridge.
Client Persistent Bytes/ sec Sent	The number of bytes contained within persistent messages sent per second by the client via the bridge.
Client Persistent Msgs Rcvd	The number of persistent messages received by the client via the bridge.
Client Persistent Msgs / sec Rcvd	The number of persistent messages received per second by the client via the bridge.
Client Persistent Msgs Sent	The number of persistent messages sent by the client via the bridge.
Client Persistent Msgs/sec Sent	The number of persistent messages sent per second by the client via the bridge.
Total Client Bytes Rcvd	The number of bytes contained within all messages received by the client via the bridge.
Total Client Bytes/sec Rcvd	The number of bytes contained within all messages received per second by the client via the bridge.
Total Client Bytes Sent	The number of bytes contained within all messages sent by the client via the bridge.
Total Client Bytes/sec Sent	The number of bytes contained within all messages sent per second by the client via the bridge.
Total Client Msgs Rcvd	The total number of all messages received by the client via the bridge.
Total Client Msgs/sec Rcvd	The total number of all messages received per second by the client via the bridge.
Total Client Msgs Sent	The total number of all messages sent by the client via the bridge.
Total Client Msgs/sec Sent	The total number of all messages sent per second by the client via the bridge.
Total Out Discards	The total number of discarded outgoing messages sent by the client via the bridge.
Total Out Discards/sec	The total number of discarded outgoing messages sent per second by the client via the bridge.
Total In Discards	The total number of discarded incoming messages received by the client via the bridge.
Total In Discards/sec	The total number of discarded incoming messages received per second by the client via the bridge.

Expired	When checked, performance data about the bridge has not been received within the time specified (in seconds) in the **$solRowExpirationTime** field in the **conf\rtvapm_solmon.properties** file. The **$solRowExpirationTimeForDelete** field allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response from the bridge. To view/edit the current values, modify the following lines in the **.properties** file:
	`# Metrics data are considered expired after this number of seconds`
	`#`
	`collector.sl.rtview.sub=$solRowExpirationTime:45`
	`collector.sl.rtview.sub=$solRowExpirationTimeForDelete:3600`
	
	In the example above, the **Expired** check box would be checked after 45 seconds, and the row would be removed from the table after 3600 seconds.

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

Bridges - Diagram

Use this topology view to monitor the health of your network bridges and VPNs. Quickly identify bridge and VPN connections, their health status and which resources their performance impacts. Drag and drop objects to arrange them on the screen (doing so does not logically impact the network bridges and VPNs). Arrows show the connections between VPNs and bridges.

Each object is a network bridge or VPN. Each is labeled with their name and color coded as follows:

- Red indicates that the object has one or more alerts in a critical state.
- Yellow indicates that the object has one or more alerts in a warning state.
- Green indicates that there are no alerts on the object.
- Gray indicates that the object is off-line.

Save: Saves the arrangement of the objects.

Restore: Returns objects to their previous positions.

Layout: Toggles between two types of layouts. One layout positions objects to the right so you might scroll in that direction to see them. The other layout pulls all the objects close together to the left, vertically, in hierarchical order.

Look at the miniature view in (upper left) to see all objects in either layout. Zoom in on an area in the topology by clicking it in the miniature view.

Drill-down to investigate in the "Bridges - Table".

To monitor network routers, VMRs and servers, see the "Neighbors - Diagram".

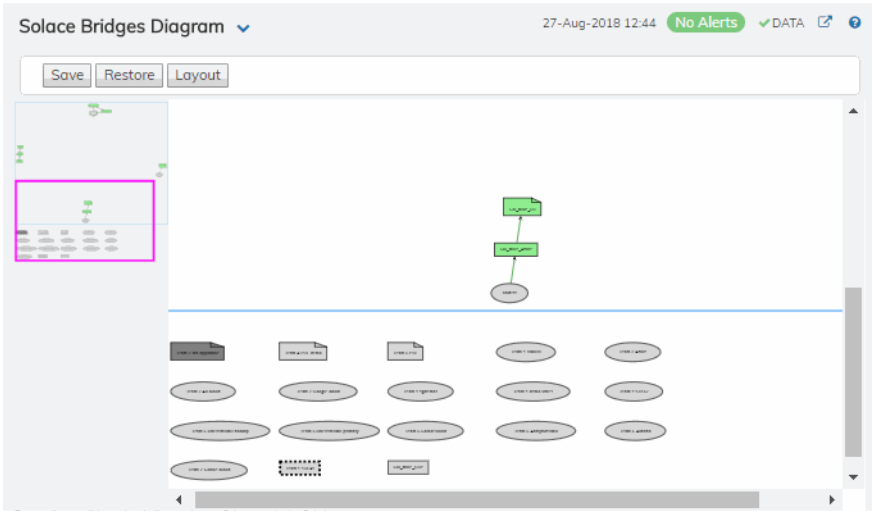

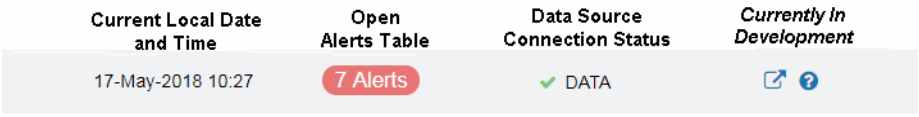

Bridge - Summary

This display allows you to view performance details for a specific bridge configured on a VPN.

Select a message router, VPN, and a bridge from the drop-down menus, and use the **Time-Range** to "zoom-in" or "zoom-out" on a specific time frame in the trend graph.

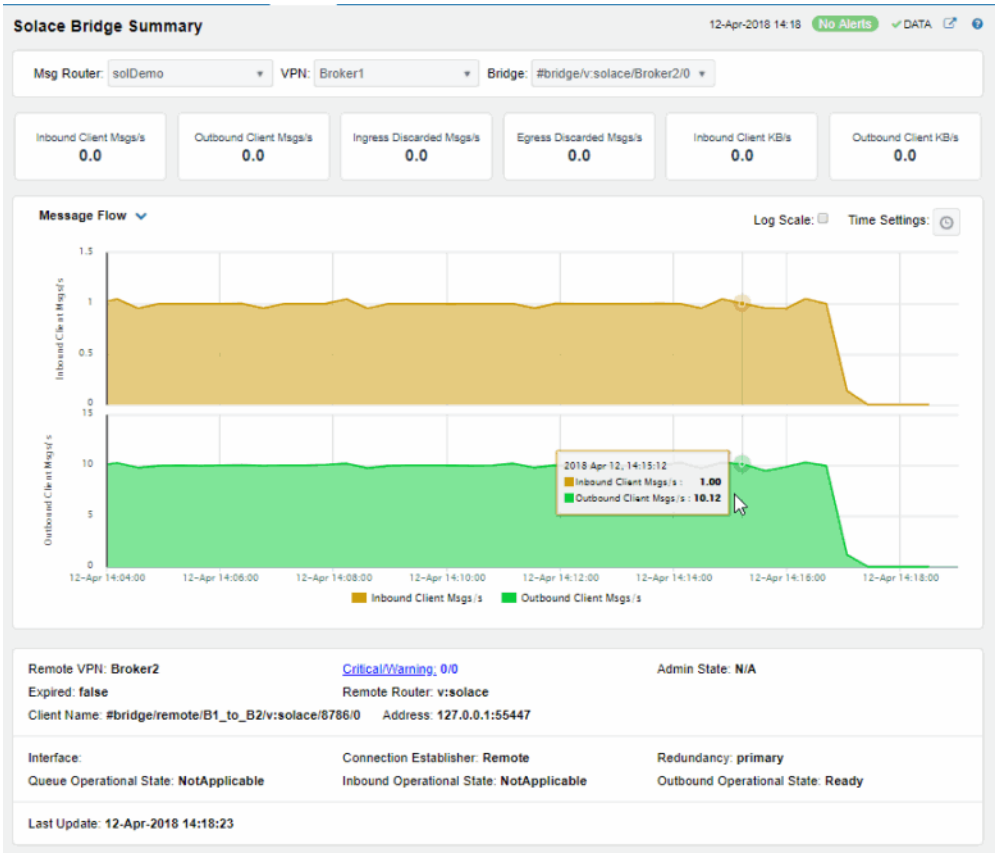

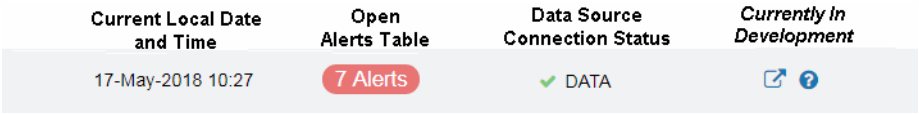

Inbound Client Msgs/s	The number of client messages received per second.
Outbound Client Msgs/s	The number of client messages sent per second.
Ingress Discarded Client Msgs/s	The number of discarded ingress messages per second.
Egress Discarded Msgs/s	The number of discarded egress messages per second.
Inbound Client KB/s	The amount of incoming client data, in KB per second.
Outbound Client KB/s	The amount of outgoing client data, in KB per second.

Messages Flow Trend Graphs
Traces the sum for the selected client.

- **Inbound Client Msgs/s**: The number of client messages received per second.
- **Outbound Client Msgs/s**: The number of client messages sent per second.

Log Scale

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

Time Settings

By default, the time range end point is the current time. To change the time range, click the **Time Settings** and either:

- choose a **Time range** from 5 Minutes to 7 Days in the drop-down menu.
- specify begin/end dates using the calendar .
- specify begin/end time using the clock .

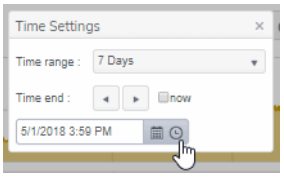

Toggle forward/backward in the trend graph per the period you choose (from the **Time range** drop-down menu) using arrows .

Restore settings to current time by selecting **now** .

Remote VPN

The name of the remote VPN that is connected to the local VPN via the bridge.

Expired

When true, performance data about the bridge has not been received within the time specified (in seconds) in the **$solRowExpirationTime** field in the **conf\rtvapm_solmon.properties** file. The **$solRowExpirationTimeForDelete** field allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response from the bridge. To view/edit the current values, modify the following lines in the **.properties** file:

```
# Metrics data are considered expired after this number
of seconds
#
collector.sl.rtview.sub=$solRowExpirationTime:45
collector.sl.rtview.sub=$solRowExpirationTimeForDelete:3
600
```

In the example above, the **Expired** check box would be checked after 45 seconds, and the row would be removed from the table after 3600 seconds.

Address

The IP address.

Interface

The interface ID.

Queue Operational State

Refer to Solace documentation for more information.

Last Update

The date and time of the last data update.

Critical/Warning

The number of critical alerts / warning alerts which also opens the **Alerts Table**.

Remote Router

The remote router.

Conn Establisher	Refer to Solace documentation for more information.
Inbound Operational State	The current inbound operational status of the bridge. (The administrator can turn off a bridge's input or output for maintenance or other reasons.)
Admin State	Indicates whether the bridge has been administratively enabled (via SolAdmin or the command line interface).
Client Name	The name of the client.
Redundancy	Indicates whether the bridge is the **primary** bridge, the **backup** bridge, the **static** bridge (default bridge used when no other bridge is available), or whether it is the only bridge available (**none**).
Outbound Op State	The current outbound operational status of the bridge. (The administrator can turn off a bridge's input or output for maintenance or other reasons.)

Endpoints

These displays list data for one or more endpoints configured on a VPN. Displays in this View are:

- "Endpoints - Table"
- "Endpoint - Summary"

Endpoints - Table

View all endpoints configured on a VPN. Each row in the table lists the details for a specific endpoint.

By default, a subset of available metrics is shown. Use **More Columns**/**Less Columns** to toggle to the complete set of metrics available (and back to the subset).

Select a router and VPN from the drop-down menus. Filter the table using the **Show Ingress Config Status Down Only** check-box ☑ and use the **Show** drop-down menus to include **All**, **Expired** or **Unexpired**.

You can click a column header to sort column data in numerical or alphabetical order, or double-click a row to drill-down and investigate in the "Endpoint - Summary" display.

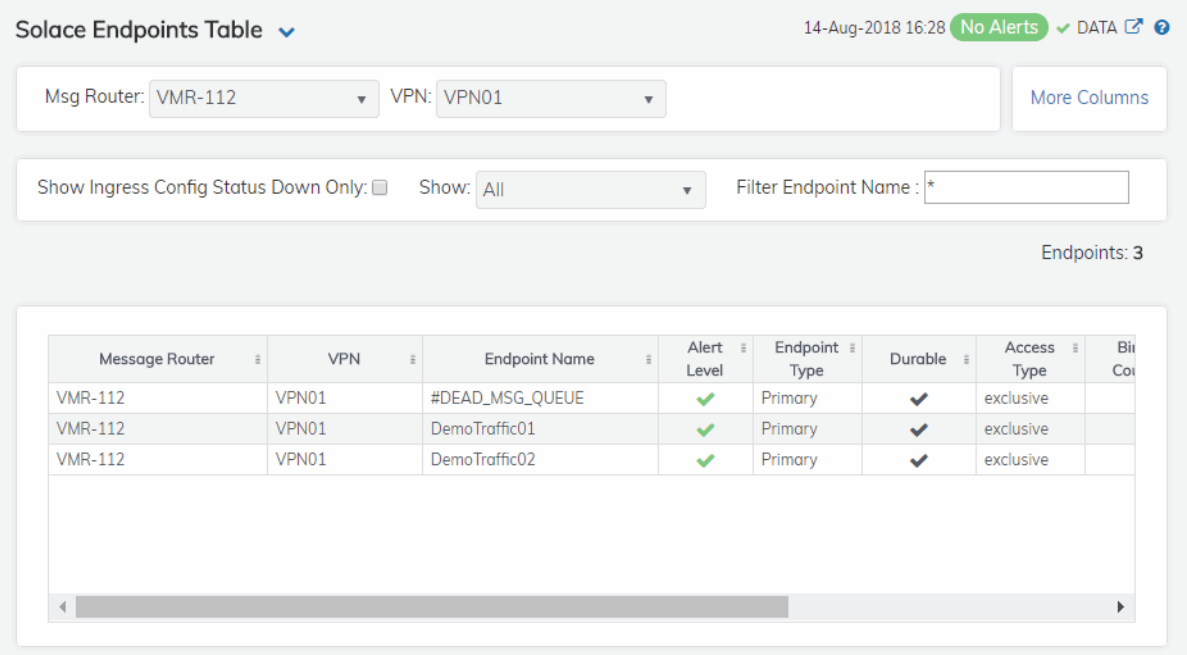

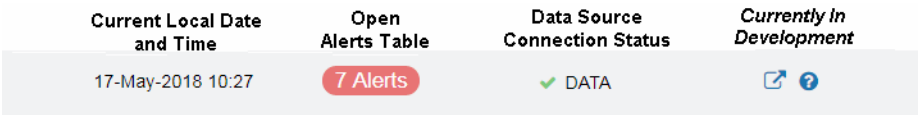

Message Router	Displays the name of the message router
VPN	The name of the VPN.
Endpoint Name	The name of the endpoint.
Alert Level	The current alert severity in the row. Red indicates that one or more metrics exceeded their ALARM LEVEL threshold. Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold. Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	The total number of active alerts for the endpoint.
Bind Count	The total number of binds connected to the endpoint.
Endpoint Type	The type of endpoint (either queue or topic).
Durable	Displays whether or not the endpoint is durable (checked) or non-durable (unchecked). Durable endpoints remain after an message router restart and are automatically restored as part of an message router's backup and restoration process.
In Config Status	Refer to Solace documentation for more information.

Out Config Status	Refer to Solace documentation for more information.
Type	Refer to Solace documentation for more information.
Access Type	Refer to Solace documentation for more information.
Pending Messages	The total number of pending messages on the endpoint.
Spool Usage (MB)	The total spool usage consumed on the endpoint (in megabytes).
High Water Mark (MB)	The highest level of spool usage on the endpoint (in megabytes).
In Selector	Refer to Solace documentation for more information.
Out Selector	Refer to Solace documentation for more information.

Expired	When checked, performance data about the endpoint has not been received within the time specified (in seconds) in the **$solRowExpirationTime** field in the **conf\rtvapm_solmon.properties** file. The **$solRowExpirationTimeForDelete** field allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response from the endpoint. To view/edit the current values, modify the following lines in the **.properties** file:
	`# Metrics data are considered expired after this number of seconds`
	`#`
	`collector.sl.rtview.sub=$solRowExpirationTime:45`
	`collector.sl.rtview.sub=$solRowExpirationTimeForDelete:3600`
	
	In the example above, the **Expired** check box would be checked after 45 seconds, and the row would be removed from the table after 3600 seconds.

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

Endpoint - Summary

This display allows you to view endpoint information, message data, and a trend graph for pending and spool messages for a specific endpoint configured on a VPN. Choose a message router, a VPN, and an endpoint from the drop-down menus, and use the **Time Settings** to "zoom-in" or "zoom-out" on a specific time frame in the trend graph.

This display is provided by default and should be used if you do not want to collect message spool data for specific VPNs. However, if you do want to configure message spool monitoring for specific VPNs, then you should use the **Single Endpoint Summary Rates** display instead, which is not included in the navigation tree by default.

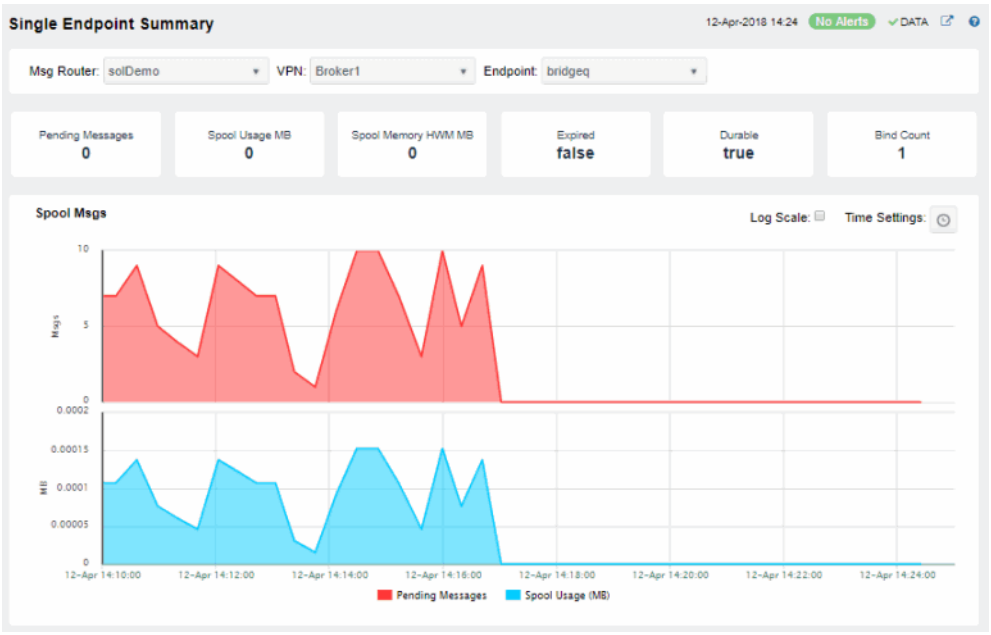

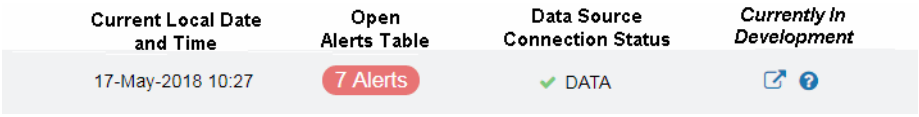

Pending Messages	The total number of pending messages on the endpoint.
Spool Usage (MB)	The current spool usage consumed on the endpoint (in megabytes).
Spool Memory HWM MB	Refer to Solace documentation for more information

Expired	When **true**, performance data about the endpoint has not been received within the time specified (in seconds) in the **$solRowExpirationTime** field in the **conf\rtvapm_solmon.properties** file. The **$solRowExpirationTimeForDelete** field allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response from the endpoint. To view/edit the current values, modify the following lines in the **.properties** file:
	`# Metrics data are considered expired after this number of seconds`
	`#`
	`collector.sl.rtview.sub=$solRowExpirationTime:45`
	`collector.sl.rtview.sub=$solRowExpirationTimeForDelete:3600`
	
	In the example above, the **Expired** check box would be checked after 45 seconds, and the row would be removed from the table after 3600 seconds.

Durable	Displays whether or not the endpoint is durable (checked) or non-durable (unchecked). Durable endpoints remain after an message router restart and are automatically restored as part of an message router's backup and restoration process.
Bind Count	The total number of binds connected to the endpoint.

Trend Graphs
Traces the sum of metrics for the endpoint.
- **Spooled Msgs**: The amount of spooled messages, in megabytes.
- **Cur Spool Usage**: The amount of space used by spooled messages, in megabytes.

Log Scale	Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
Base at Zero	Select to use zero (**0**) as the Y axis minimum for all graph traces.
Time Settings	By default, the time range end point is the current time. To change the time range, click the **Time Settings** and either: • choose a **Time range** from 5 Minutes to 7 Days in the drop-down menu. • specify begin/end dates using the calendar . • specify begin/end time using the clock . 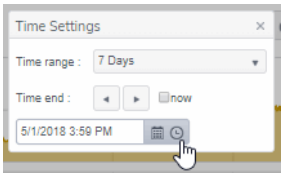 Toggle forward/backward in the trend graph per the period you choose (from the **Time range** drop-down menu) using arrows . Restore settings to current time by selecting **now** .

Endpoint Type	The type of endpoint.
Egress Config Status	Refer to Solace documentation for more information.
Egress Selector Present	Refer to Solace documentation for more information.
Last Update	The date and time of the last data update.
Critical/Warning	The number of critical alerts / warning alerts which also opens the **Alerts Table**.
Access Type	Refer to Solace documentation for more information.
Ingress Config Status	Refer to Solace documentation for more information.
Ingress Selector Present	Refer to Solace documentation for more information.

Capacity

These displays provide current router capacity metrics, alert count and severity at the message router level. Displays in this View are:

- "Capacity - Table": View client, spool usage, incoming messages, outgoing messages, incoming bytes, and outgoing bytes data for all message routers.
- "Capacity - Summary": View client, spool usage, incoming messages, outgoing messages, incoming bytes, and outgoing bytes data for a specific message router.
- "Capacity - Trends": View the message router capacity data for a specific message router in a trend graph format.

Capacity - Table

View current and HWM (high water mark for the last 30 days) capacity utilization data for all message routers.

By default, a subset of available metrics is shown. Use **More Columns**/**Less Columns** to toggle to the complete set of metrics available (and back to the subset).

You can view client, spool usage, incoming message, outgoing message, incoming bytes, and outgoing bytes data for the message router. Each table row is a different message router.

Double-click a row to drill-down and investigate in the "Capacity - Summary" display.

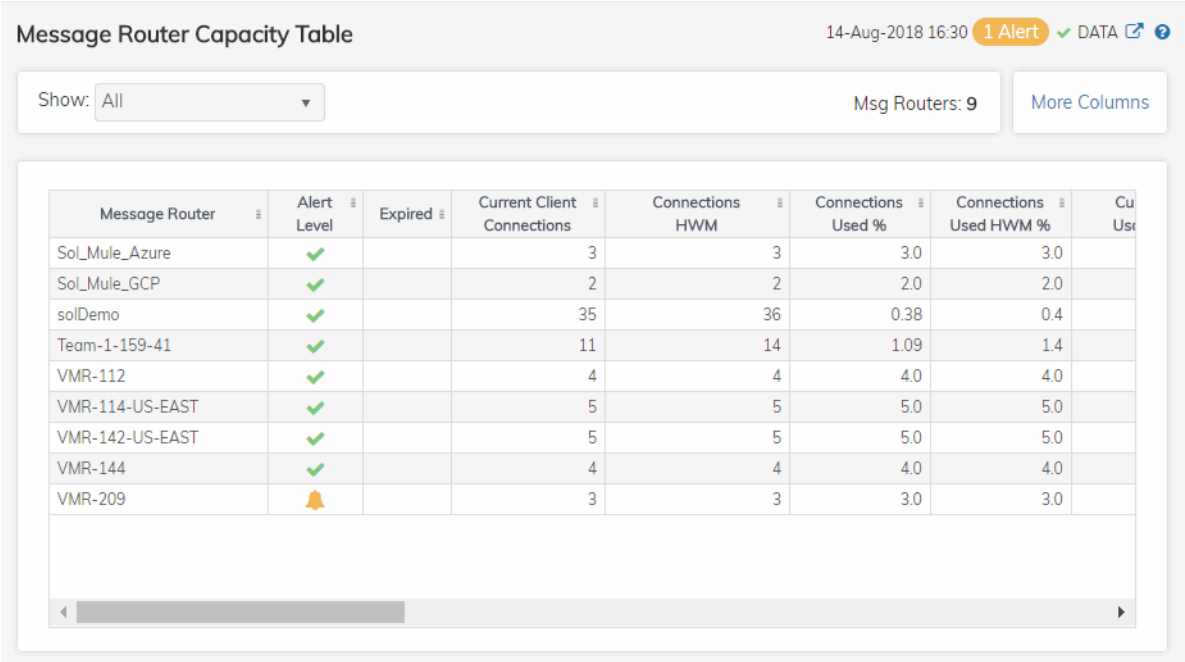

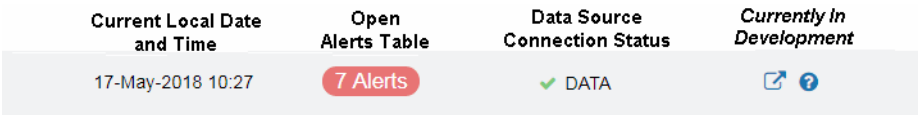

Message Router	The name of the message router.
Alert Level	The maximum level of alerts in the row: Red indicates that one or more metrics exceeded their ALARM LEVEL threshold. Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold. Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	The total number of active alerts.
Current Client Connections	The current number of clients connected.
Connections HWM	The greatest number of connections in the last 30 days.
Connections Max	The greatest number of connections since the message router last started.
Connections Reserved	The current number of reserved connections.
Connections Used %	The current amount of connections used, in percent.
Connections Used HWM %	The greatest amount of connections used, in percent, in the last 30 days.
Cur Spool Usage MB	The current amount of used spool disk, in megabytes.
Cur Spool Usage HWM	The greatest amount of spool disk used in the last 30 days.
Spool Disk Allocated	The amount of allocated spool disk.
Spool Reserved	The amount of reserved spool disk.
Current Spool Usage %	The current amount of used spool disk, in percent.
Current Spool Usage % HWM	The greatest amount of used spool disk in the last 30 days, in percent.
Delivered Unacked Msgs Util %	Refer to Solace documentation for more information.
Ingress Flow Count	The number of ingress flows.
Ingress Flow HWM	The greatest number of ingress flows in the last 30 days.
Ingress Flows Allowed	The maximum number of ingress flows allowed.
Ingress Flow Count %	The amount of ingress flows in percent.
Ingress Flow Count HWM %	The greatest amount of ingress flows in the last 30 days, in percent.
Ingress Msgs/s	The number of ingress messages per second.
Ingress Msgs/s HWM	The greatest number of ingress messages per second in the last 30 days.
Max Ingress Msgs/s	The maximum number of ingress flows per second allowed.
Ingress Msgs %	The amount of ingress messages in percent.

Ingress Msgs/s HWM %	The greatest amount of ingress messages in the last 30 days, in percent.
Cur Egress Msgs/s	The number of egress messages per second.
Egress Msgs/s HWM	The greatest number of egress messages per second in the last 30 days.
Max Egress Msgs/s	The maximum number of egress flows per second allowed.
Egress Msgs %	The amount of egress messages in percent.
Egress Msgs/s HWM %	The greatest amount of ingress messages in the last 30 days, in percent.
Cur Egress Bytes/s	The amount of egress in bytes per second.
Egress Bytes/s HWM	The greatest amount of egress, in bytes per second, in the last 30 days, in percent.

Expired	When checked, performance data about the VPN has not been received within the time specified (in seconds) in the **$solRowExpirationTime** field in the **conf\rtvapm_solmon.properties** file. The **$solRowExpirationTimeForDelete** field allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response from the VPN. To view/edit the current values, modify the following lines in the **.properties** file:
	`# Metrics data are considered expired after this number of seconds`
	`#`
	`collector.sl.rtview.sub=$solRowExpirationTime:45`
	`collector.sl.rtview.sub=$solRowExpirationTimeForDelete:3600`
	
	In the example above, the **Expired** check box would be checked after 45 seconds, and the row would be removed from the table after 3600 seconds.

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

Capacity - Summary

This display, a pivoted view of the "Capacity - Table", allows you to view current and HWM (high water mark for the last 30 days) capacity utilization data for a single message router. Select a router from the drop-down menu to view client, spool usage, incoming message, outgoing message, incoming bytes, and outgoing bytes data for the message router.

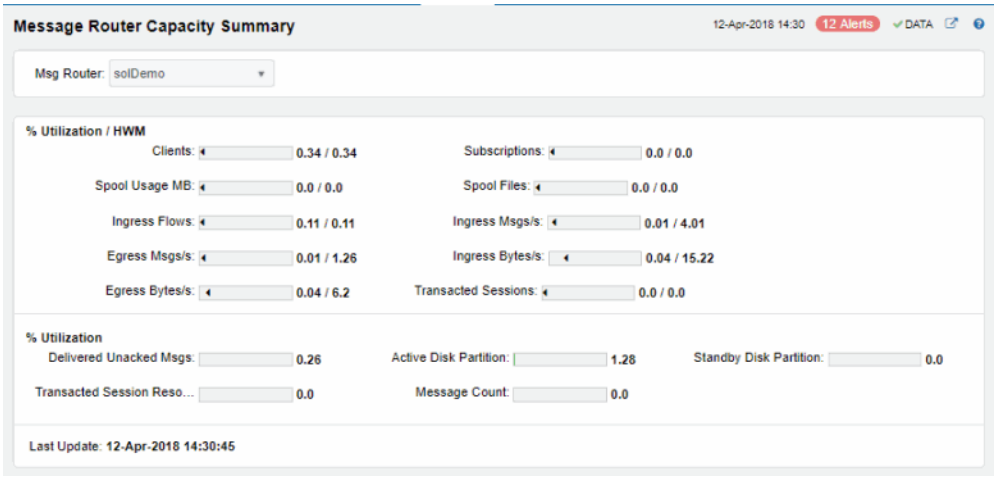

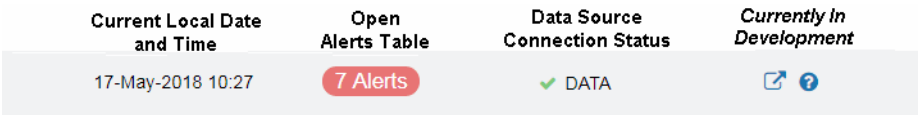

% Utilization/HWM These values show high water marks (peak capacity utilization) for the last 30 days.

Clients	The current number of clients connected to the message router.
Spool Files	The highest number of spool files on the message router in the past 30 days.
Egress Msgs/s	The highest number of outgoing messages per second on the router in the past 30 days.
Transacted Sessions	The highest number of transacted sessions on the message router in the last 30 days.
Subscriptions	The highest number of subscriptions on the message router in the last 30 days.
Ingress Flows	The highest number of inflows on the message router in the last 30 days.
Ingress Bytes/s	The highest amount of inflows, in bytes per second, on the router in the past 30 days.
Spool Usage MB	The highest amount of spool utilization, in megabytes per second, on the router in the past 30 days.
Ingress Msgs/s	The highest number of incoming messages per second on the router in the past 30 days.
Egress Bytes/s	The highest number of outgoing messages per second on the router in the past 30 days.

% Utilization These values show current capacity utilization.

Delivered Unacked Msgs	The current number of delivered messages that were not acknowledged divided by the maximum number of delivered messages that were not acknowledged allowed on the message router.

Transacted Sessions Reso...	The current number of transacted sessions that were resolved on the message router.
Active Disk Partition	The percentage of available active disk partition that is used.
Message Count	The current number of messages on the message router.
Standby Disk Partition	The percentage of available standby disk partition that has been used.
Last Update	The date and time of the last data update.

Capacity - Trends

This display allows you to view a trend graph that traces router performance data for clients & spool data, message flow and throughput. Select a router and a performance metric from the drop-down menus.

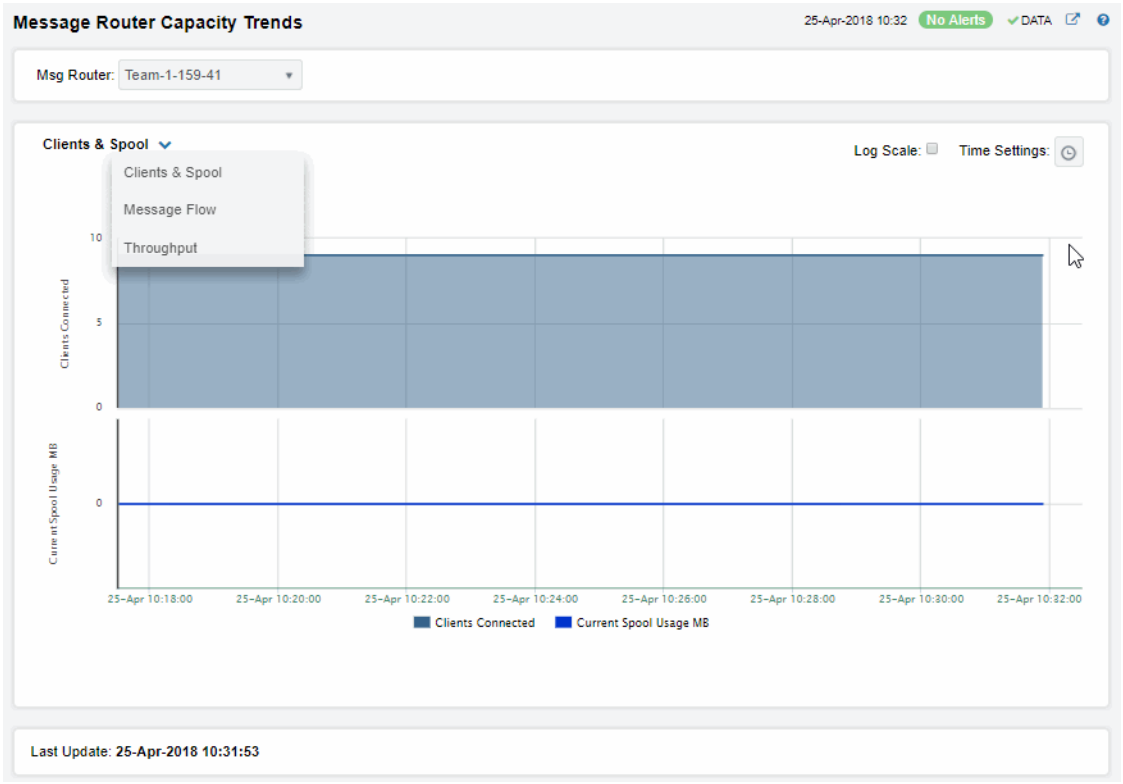

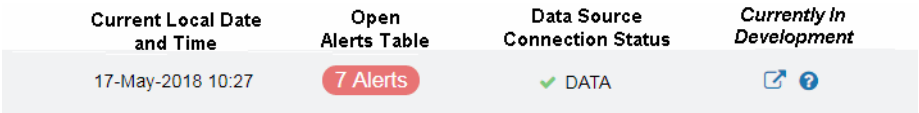

Clients & Spool	The trend graph traces the following performance metrics: **Clients Connected**: The current number of clients connected to the message router. **Current Spool Usage**: The current spool usage, in megabytes, on the message router.
Message Flow	The trend graph traces the following: **Ingress Msgs/sec**: The number of incoming messages per second on the router. **Egress Msgs/sec**: The number of outgoing messages per second on the router.
Throughput	The trend graph traces the following: **Ingress KB/sec**: The amount of incoming per second, in KB, on the router. **Egress KB/sec**: The number of outgoing data per second, in KB, on the router.
Log Scale	Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
Base at Zero	Select to use zero (**0**) as the Y axis minimum for all graph traces.
Time Settings	By default, the time range end point is the current time. To change the time range, click the **Time Settings** and either: • choose a **Time range** from 5 Minutes to 7 Days in the drop-down menu. • specify begin/end dates using the calendar . • specify begin/end time using the clock . 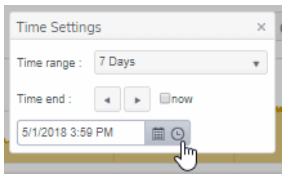 Toggle forward/backward in the trend graph per the period you choose (from the **Time range** drop-down menu) using arrows . Restore settings to current time by selecting **now** .

Alerts Table

Use this display to track and manage all Solace 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
	The alert has been resolved. An alert is automatically cleared when the value being monitored no longer in the alert threshold.

Use No Alerts in the title bar to access the **Alerts Table**.

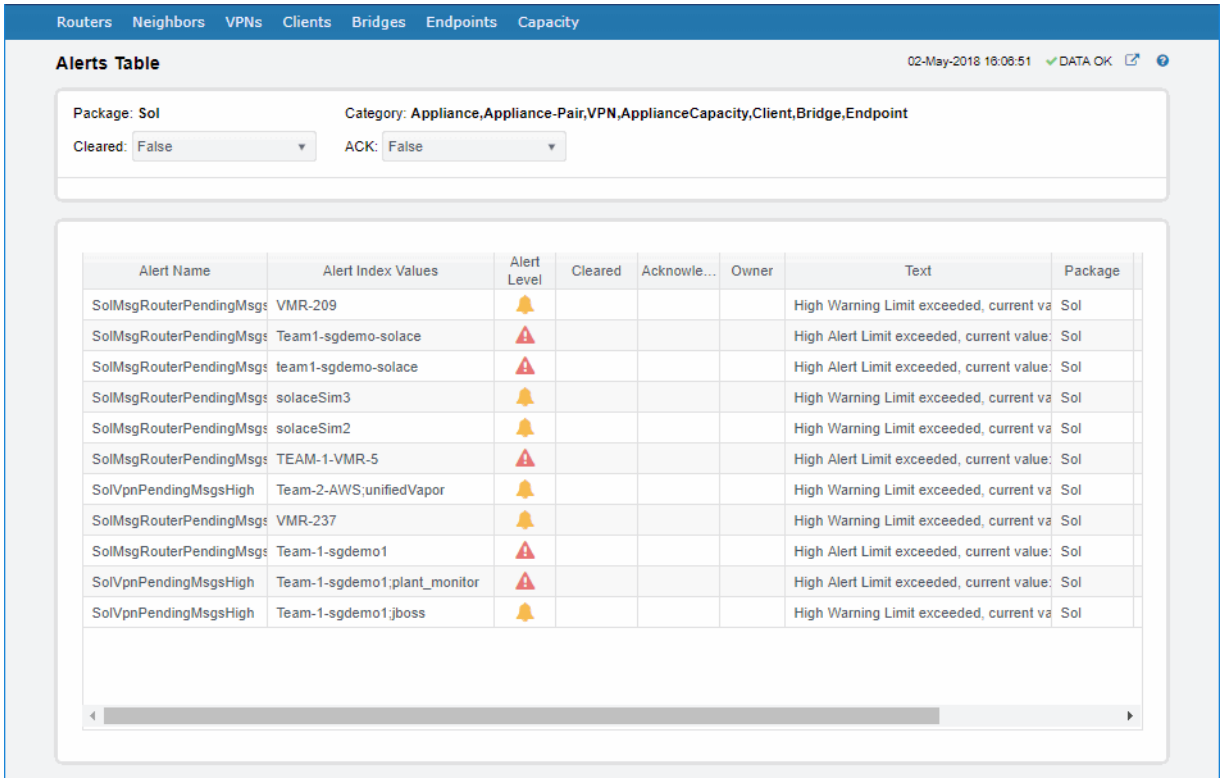

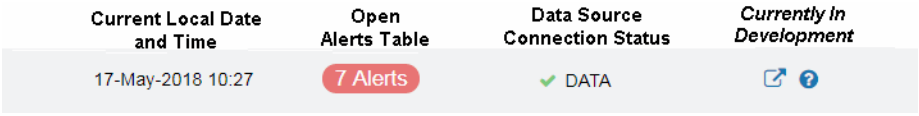

Alert Count shows the current number of alerts in the table. Filter the alert list using the drop-down menus where:

Cleared:
- **All** - See cleared and uncleared alerts.
- **True** - See ONLY cleared alerts.
- **False** - See ONLY uncleared alerts.

ACK:
- **All** - See acknowledged and unacknowledged alerts.
- **True** - See ONLY acknowledged alerts.
- **False** - See ONLY unacknowledged alerts.

Alert Name The name of the alert. For a complete list of available alerts, see the Alert Administration display.

Alert Index Values The IP address and port number for the source (application, server, and so forth) associated with the alert.

Alert Level	The maximum level of alerts in the row: Red indicates that one or more metrics exceeded their ALARM LEVEL threshold. Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold. Green indicates that no metrics have exceeded their alert thresholds.
Cleared	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.
Acknowledged	When checked, this typically indicates that the alert is being addressed.
Owner	The named owner assigned by the administrator.
Text	Descriptive text about the alert.
Package	The type of technology associated with the alert (for example, Solace).
Category	The type of technology component associated with the alert (for example, Endpoint).
Row Update Time	The date and time that the row was last updated.

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:

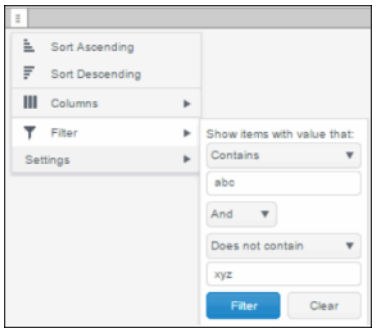

Options in the **Column Filter** dialog vary according to the data type of the selected column:

- **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.

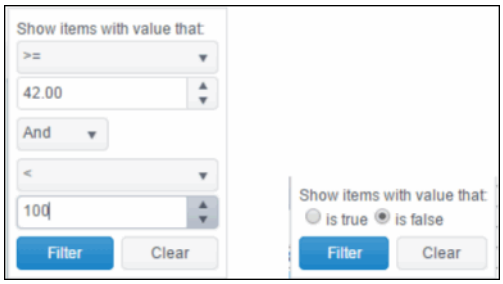

- **Date columns**: You can select a date and time and choose whether matching items should have a timestamp that is the same as, before, or after the filter time. The date is selected by clicking on the calendar icon and picking a date from a calendar dialog. The time is selected by clicking on the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

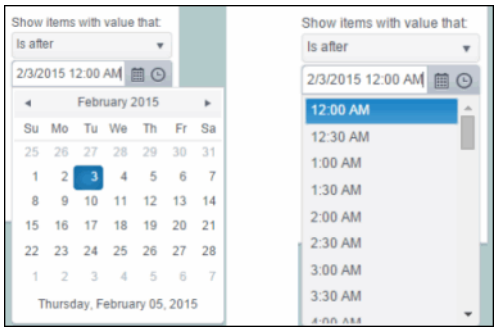

  Alternatively, a date and time can be typed into the edit box. The strings shown in a date column are formatted by the Display Server using its time zone. But if a filter is specified on a date column, the date and time for the filter are computed using the client system's time zone. This can be confusing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

For more information about table features, see the complete RTView Monitor for Solace User's Guide.

Administration

These displays enable you to set alert thresholds, observe how alerts are managed, and view internal data gathered and stored by RTView (used for troubleshooting with SL Technical Support). Displays in this View are:

- "Alert Administration": Displays active alerts and provides interface to modify and manage alerts.
- "Alert Administration Audit": View cached data that RTView is capturing and maintaining, and use this data use this for debugging with SL Technical Support.
- "RTView Cache Tables": Display information about RTView Agent data servers.
- "RTView Agent Admin": Display information about RTView Agent data servers.

Alert Administration

This section includes:

- "Tabular Alert Administration" on page 126
- "Setting Override Alerts" on page 143

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.

For details about alerts for Solace, see **Appendix A, Alert Definitions**.

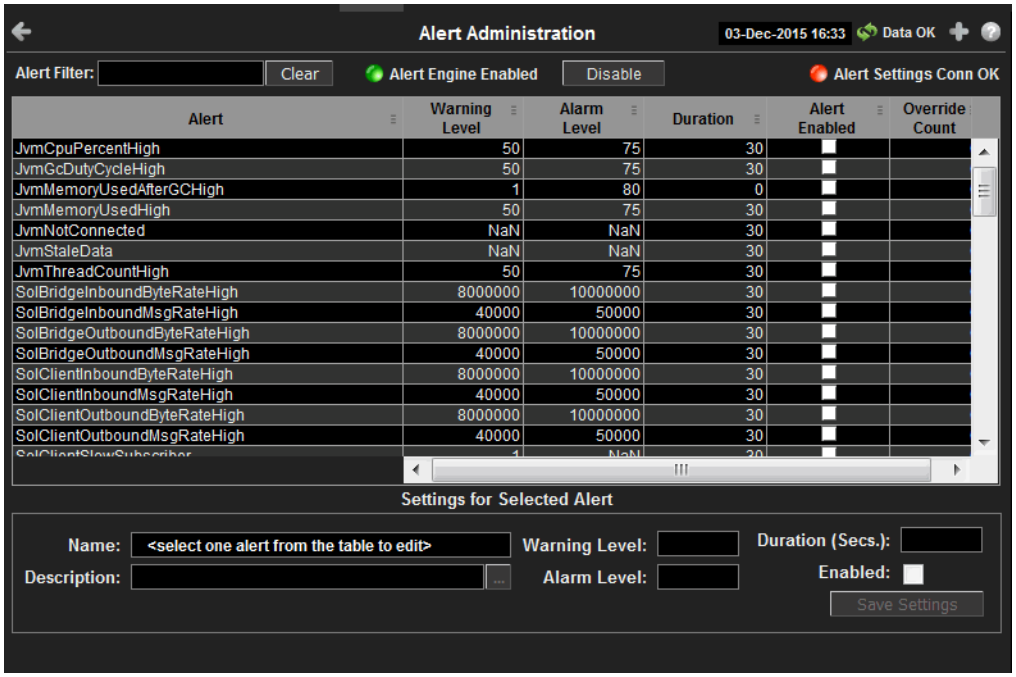

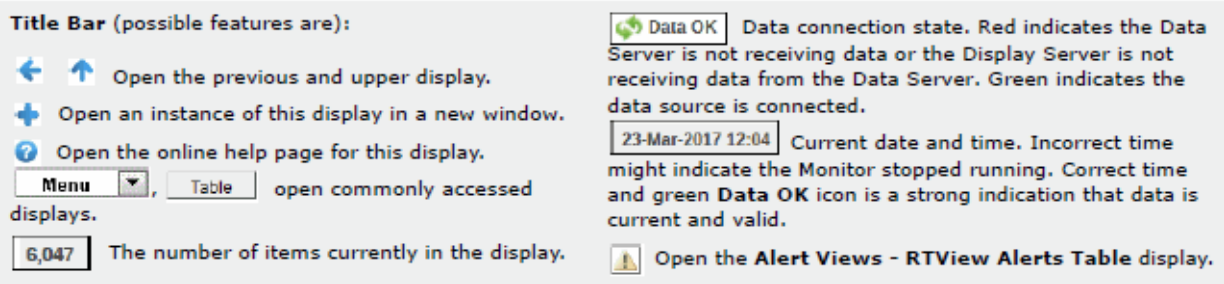

Fields and Data
This display includes:

Alert Filter Enter the (case-sensitive) string to filter the table by the **Alert** table column value. NOTE: Partial strings can be used without wildcard characters. Press **<enter>** or click elsewhere in the display to apply the filter.

Clear Clears the **Alert Filter** entry.

Alert Settings The Alert Server connection state:
- Disconnected.
- Connected.

Active Alert Table
This table describes the global settings for all alerts on the system. Select an alert. The name of the selected alert populates the **Settings for Selected Alert Name** field (in the lower panel). Edit **Settings for Selected Alert** fields and click **Save Settings** when finished.

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. A value of: **-0** indicates that no overrides are applied to the alert. **-1** indicates that the alert does not support overrides.

Settings for Selected Alert
To view or edit Global settings, select an alert from the **Active Alert Table**. Edit the **Settings for Selected Alert** fields and click **Save Settings** when finished.

To set override alerts, click on **Override Settings** to open the **Tabular Alert Administration** display.

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**.

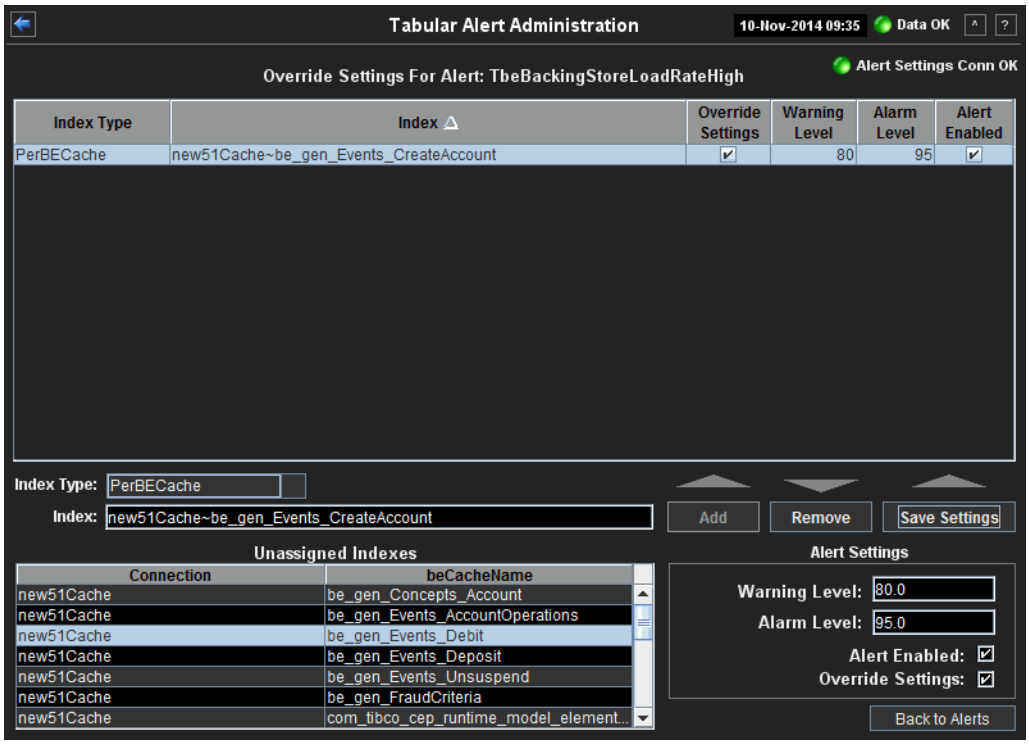

Fields and Data
This display includes:

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.

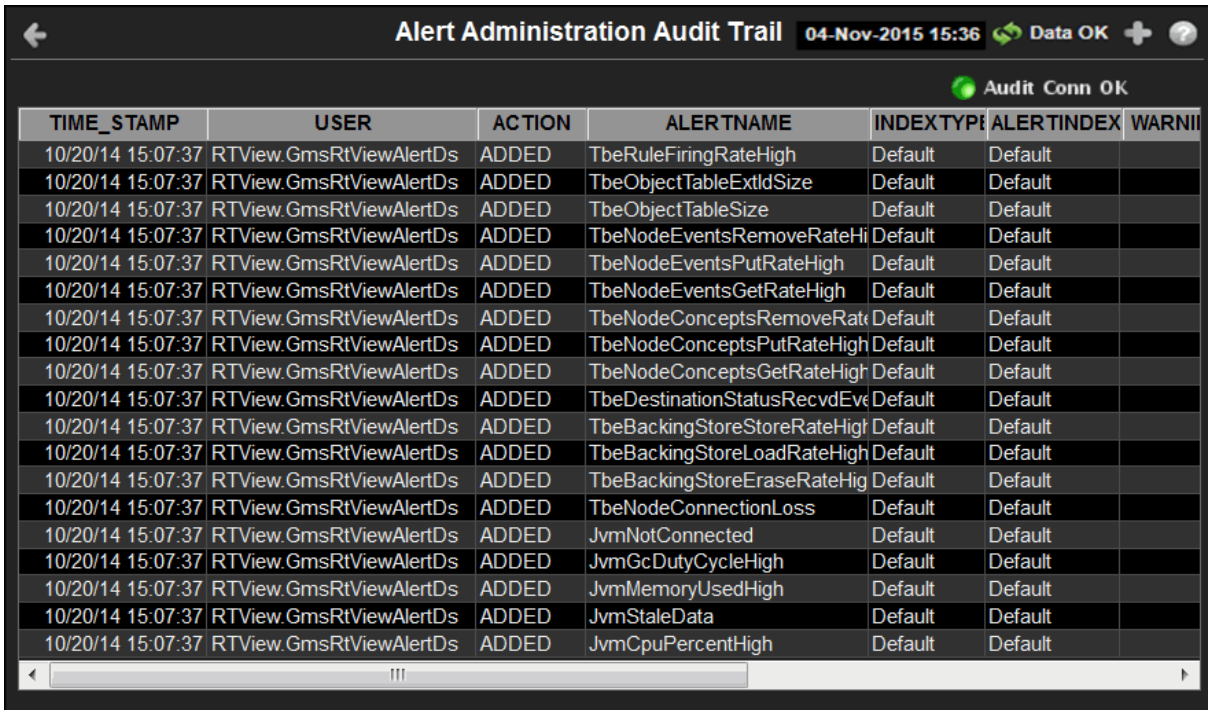

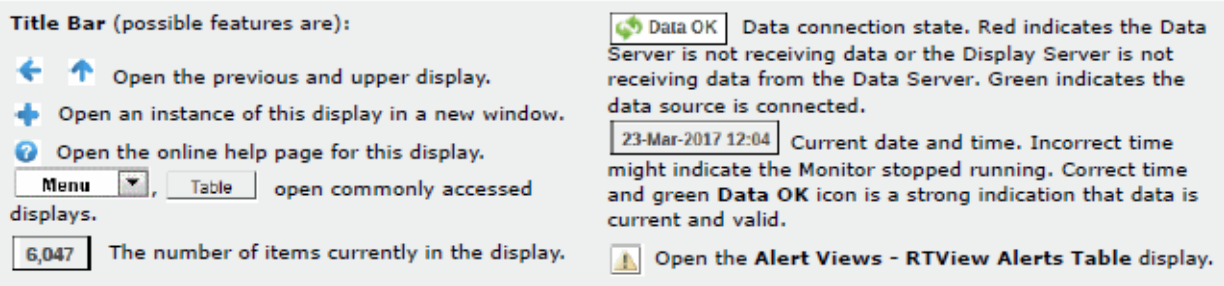

Audit Conn OK	The Alert Server connection state: Disconnected. Connected.
TIME_STAMP	The date and time of the modification.
USER	The user name of the administrator who made the modification.
ACTION	The type of modification made to the alert, such as UPDATED.
ALERTNAME	The name of the alert modified.
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).

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.

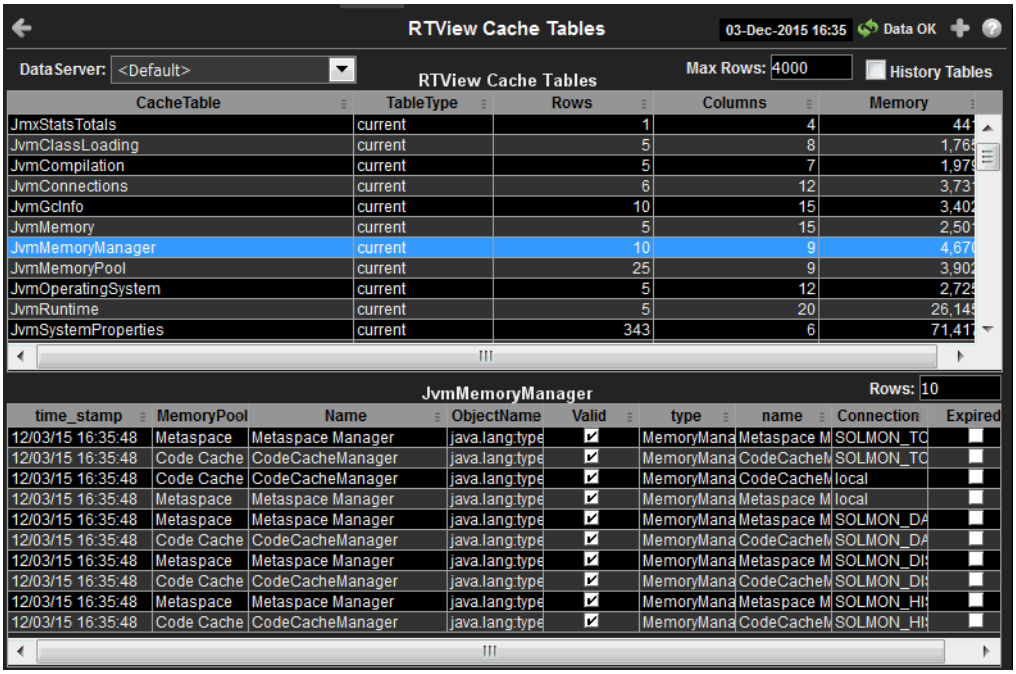

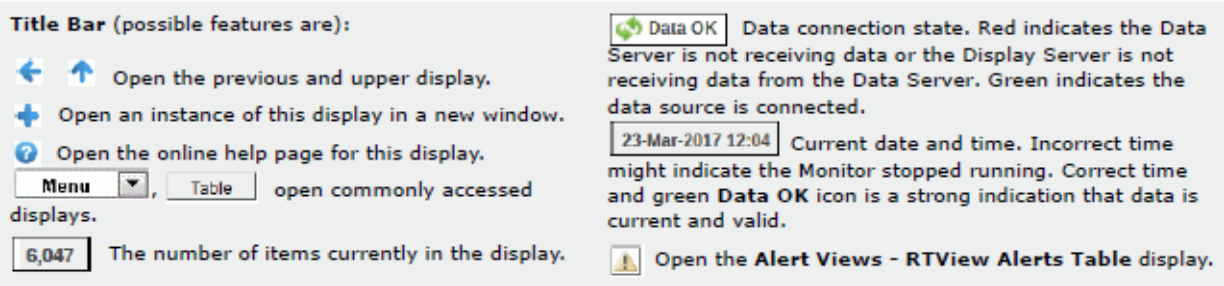

DataServer Select a data server from the drop down menu.

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.

RTView Agent Admin

Verify when agent metrics were last queried by the Monitor. The data in this display is predominantly used for debugging by Technical Support.

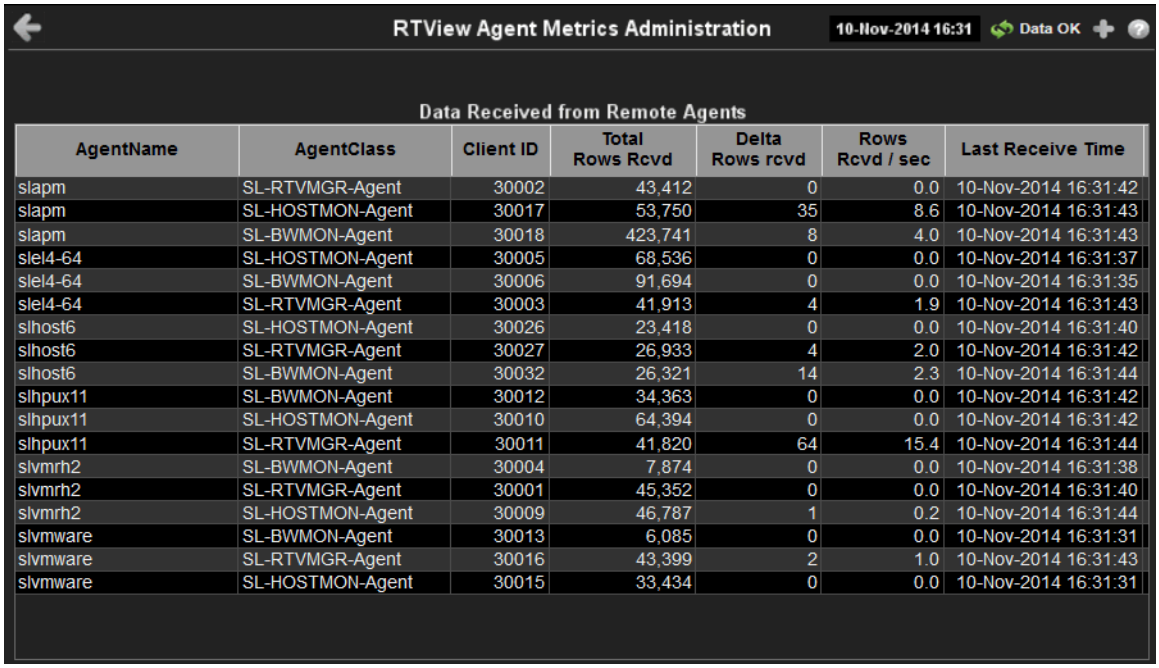

AgentName	AgentClass	Client ID	Total Rows Rcvd	Delta Rows rcvd	Rows Rcvd / sec	Last Receive Time
slapm	SL-RTVMGR-Agent	30002	43,412	0	0.0	10-Nov-2014 16:31:42
slapm	SL-HOSTMON-Agent	30017	53,750	35	8.6	10-Nov-2014 16:31:43
slapm	SL-BWMON-Agent	30018	423,741	8	4.0	10-Nov-2014 16:31:43
slel4-64	SL-HOSTMON-Agent	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.

RTView Manager for Solace Displays

Use the RTView Manager for Solace to set the following for Solace components you are monitoring: alert thresholds, manage alerts, caches, remote agents and Syslog events. The **Alert Administration** display opens by default:

The RTView Manager for Solace has the following Views:

- "Syslog"
- "Alert Views"
- "Administration"

Syslog

The display in this View provides a tabular list of all Syslog events:

- "All Syslog Events Table" on page 134: View all Syslog events for all your Solace message routers.

All Syslog Events Table

This table lists all Syslog events collected from one or all Solace message routers. Each row in the table is a different message. Filter messages per single Solace message router or all message routers (choose **All Hosts** from the **Source** drop-down menu), a single tag or **All Tags**, a single severity level or all levels (choose **All Levels** from the **Severity** drop-down menu), and specify a **Time Settings**.

Click a column header to sort column data in numerical, alphabetical or chronological order.

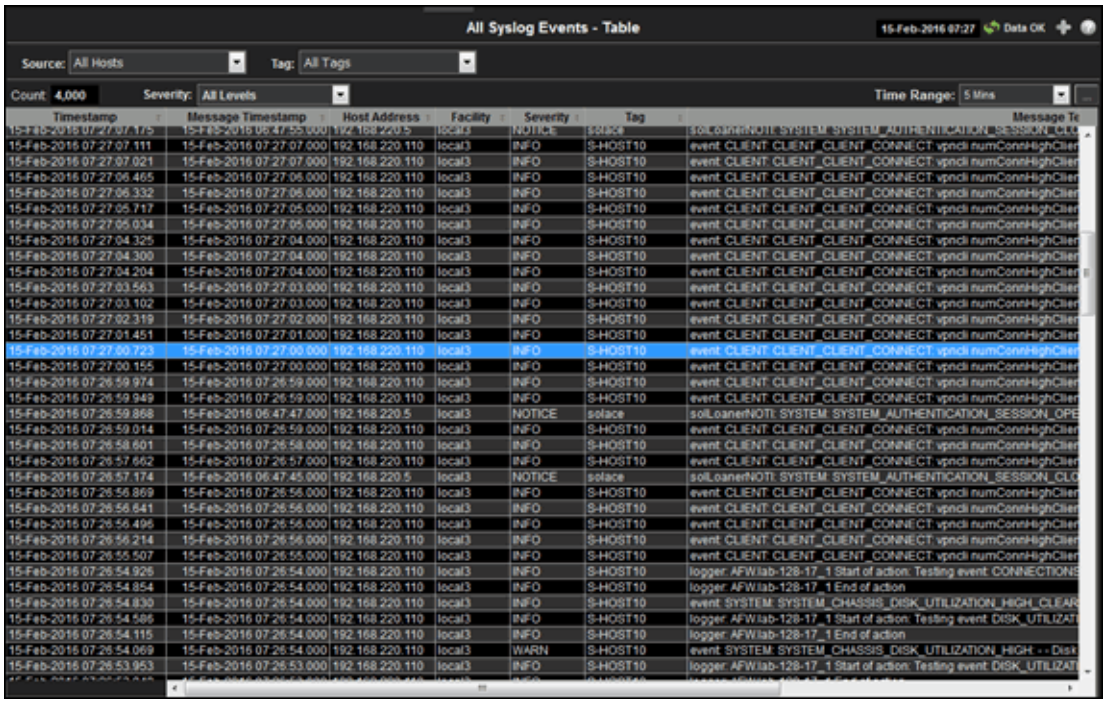

Title Bar (possible features are):

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.

Source: Select the host for which you want to view data, or **All Hosts**.

Tag: Select the message tag for which you want to view data, or **All Tags**.

Severity: Select the message severity level for which you want to view data, or **All Levels**.

Time Settings: By default, the time range end point is the current time. To change the time range, click the **Time Settings** and either:

- choose a **Time range** from 5 Minutes to 7 Days in the drop-down menu.
- specify begin/end dates using the calendar .
- specify begin/end time using the clock .

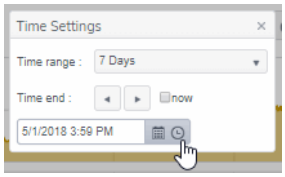

Toggle forward/backward in the trend graph per the period you choose (from the **Time range** drop-down menu) using arrows .

Restore settings to current time by selecting **now** .

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

Message Timestamp The date and time the message was sent.

Host Address The host IP address. Refer to Solace documentation for more information.

Facility The message facility code. Refer to Solace documentation for more information.

Severity The message severity level. Refer to Solace documentation for more information.

- **INFO**
- **NOTICE**
- **NOTICE or higher**
- **WARN**
- **WARN or higher**
- **ERROR**
- **ERROR or higher**
- **CRITICAL**
- **ALERT**
- **EMERGENCY**

Tag The host name. Refer to Solace documentation for more information.

Message Text The content of the message.

Alert Views

This display presents detailed information about all Solace component alerts that have occurred in your monitoring system.

Alert Detail Table

Use this display to track and manage all Solace component 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.

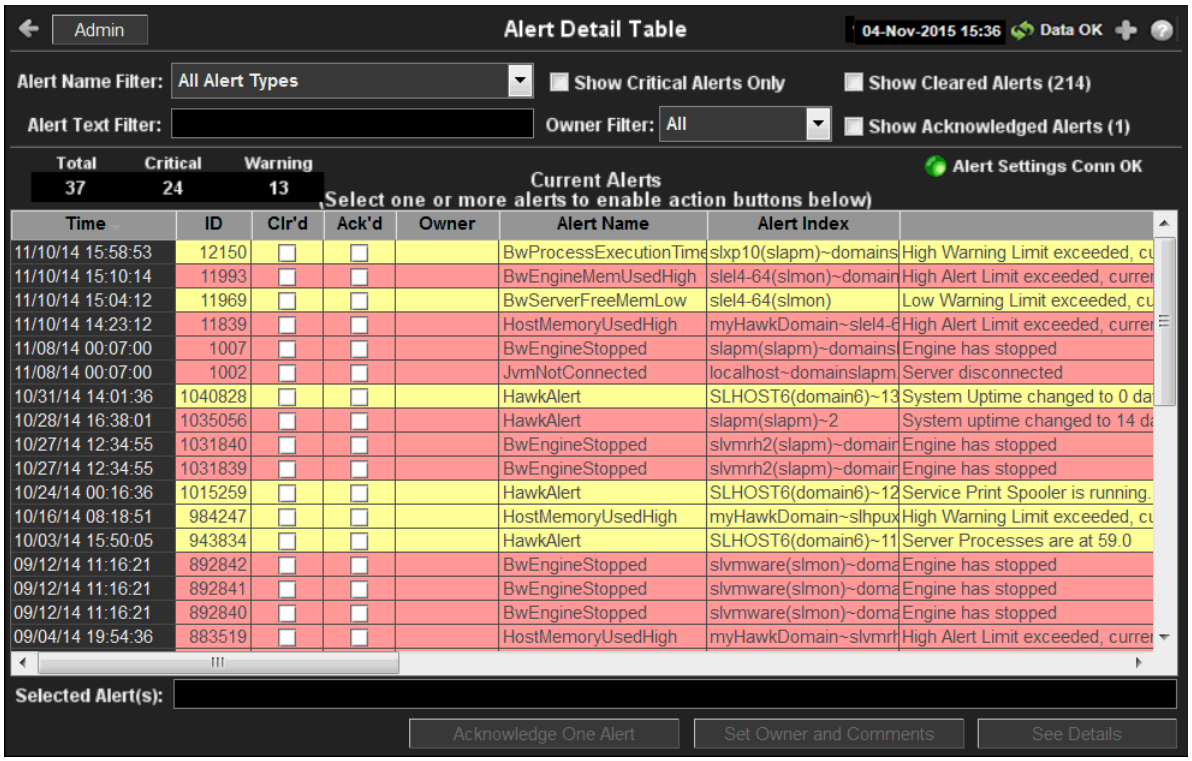

Title Bar (possible features are):

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.

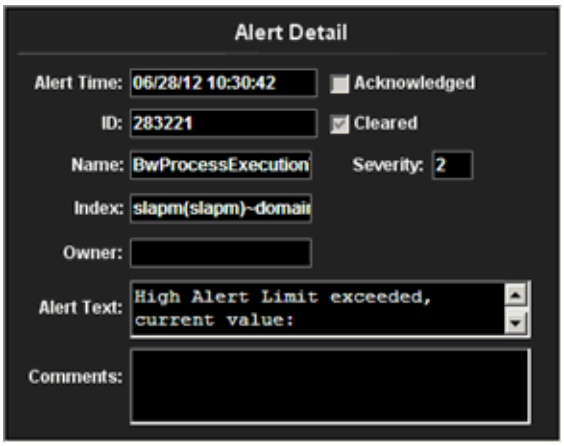

Administration

These displays enable you to set alert thresholds for Solace components, observe how alerts are managed, and view internal data gathered and stored by RTView (used for troubleshooting with SL Technical Support). Displays in this View are:

- "Alert Administration": Displays active alerts and provides interface to modify and manage alerts.
- "Alert Administration Audit"
- "RTView Cache Tables": View cached data that RTView is capturing and maintaining, and use this data use this for debugging with SL Technical Support.
- "RTView Agent Administration": Display information about RTView Agent data servers.
- "About"

Alert Administration

This section includes:

- "Global Thresholds"
- "Override Thresholds"
- "Tabular 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.

For details about alerts for Solace, seethe **Appendix for Alert Definitions**.

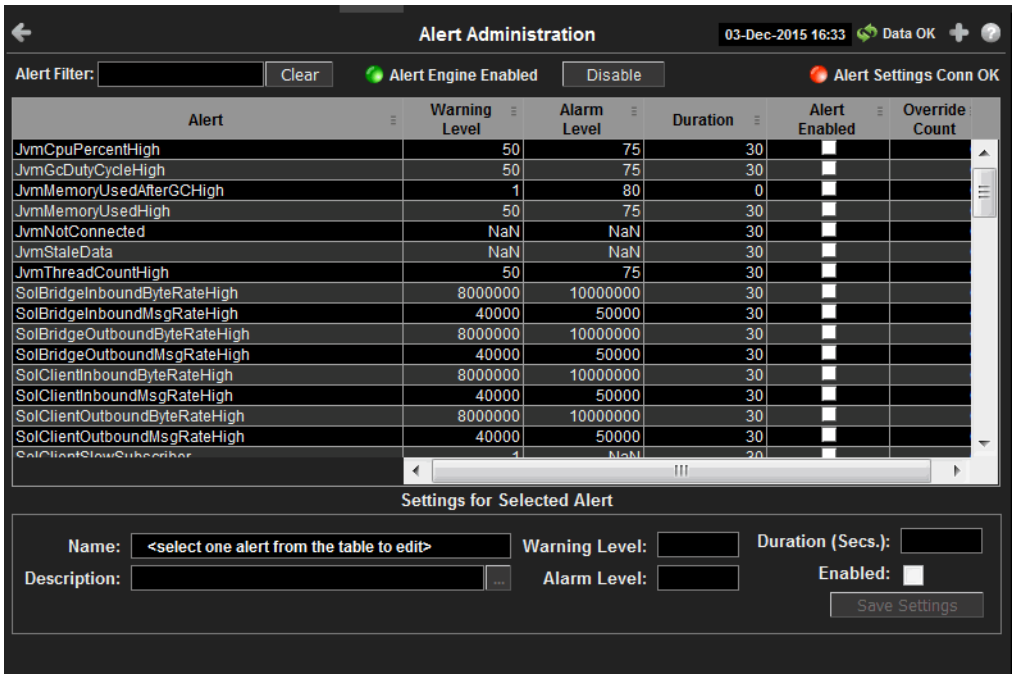

Title Bar (possible features are):

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.

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**.

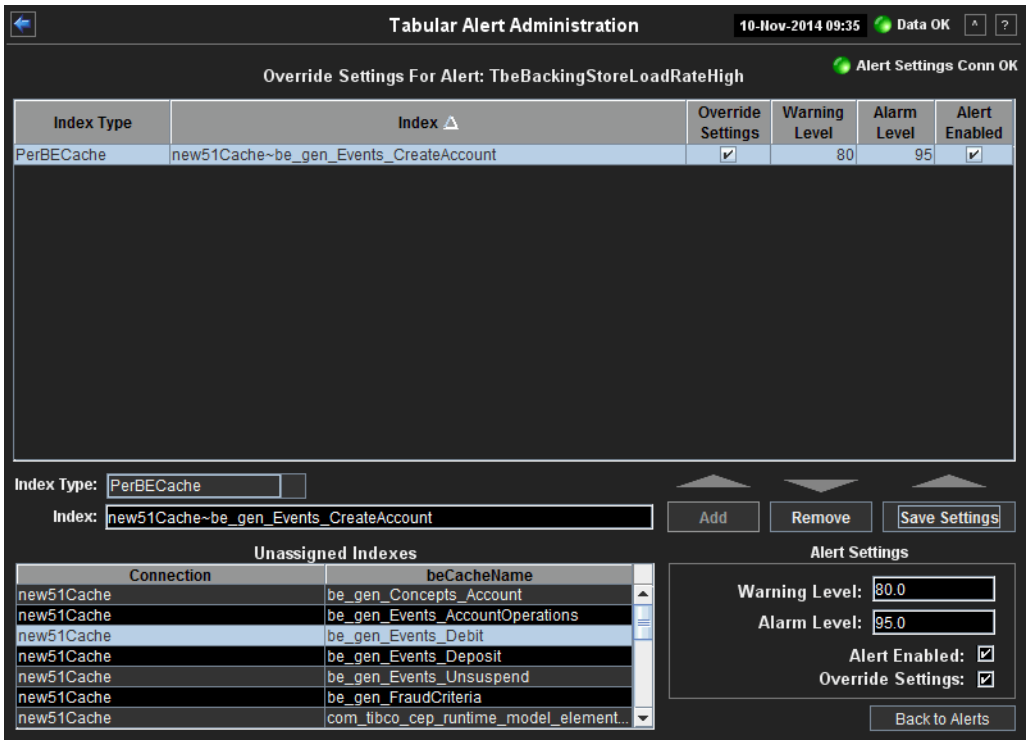

Fields and Data
This display includes:

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, click Sort to order the **ALERTNAME** column.

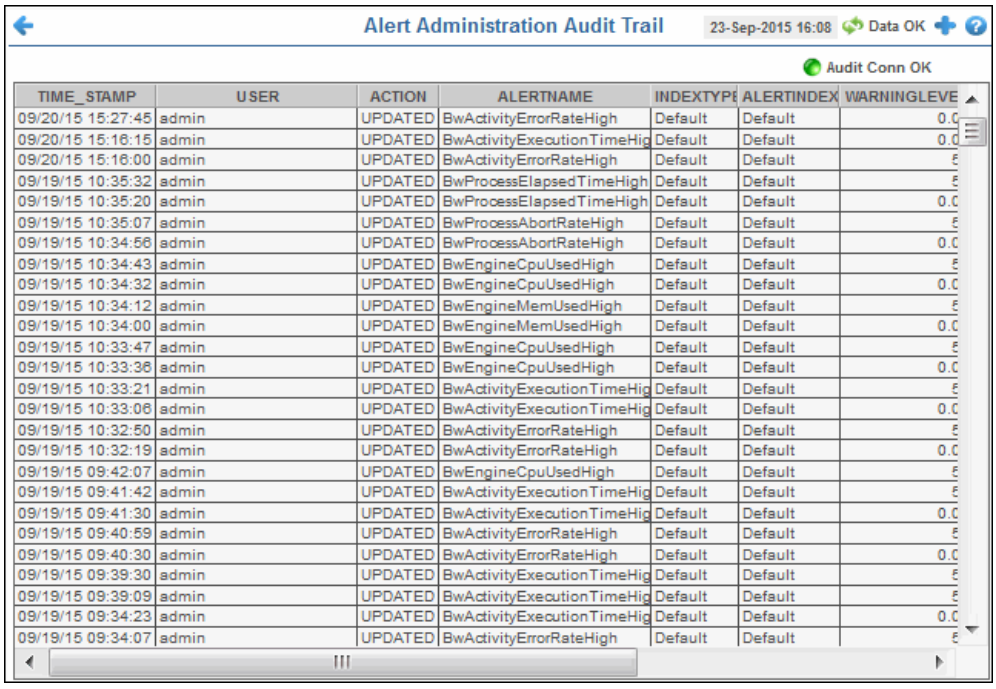

TIME_STAMP	USER	ACTION	ALERTNAME	INDEXTYPE	ALERTINDEX	WARNINGLEVE
09/20/15 15:27:45	admin	UPDATED	BwActivityErrorRateHigh	Default	Default	0.0
09/20/15 15:16:15	admin	UPDATED	BwActivityExecutionTimeHig	Default	Default	0.0
09/20/15 15:16:00	admin	UPDATED	BwActivityErrorRateHigh	Default	Default	5
09/19/15 10:35:32	admin	UPDATED	BwProcessElapsedTimeHigh	Default	Default	5
09/19/15 10:35:20	admin	UPDATED	BwProcessElapsedTimeHigh	Default	Default	0.0
09/19/15 10:35:07	admin	UPDATED	BwProcessAbortRateHigh	Default	Default	5
09/19/15 10:34:56	admin	UPDATED	BwProcessAbortRateHigh	Default	Default	0.0
09/19/15 10:34:43	admin	UPDATED	BwEngineCpuUsedHigh	Default	Default	5
09/19/15 10:34:32	admin	UPDATED	BwEngineCpuUsedHigh	Default	Default	0.0
09/19/15 10:34:12	admin	UPDATED	BwEngineMemUsedHigh	Default	Default	5
09/19/15 10:34:00	admin	UPDATED	BwEngineMemUsedHigh	Default	Default	0.0
09/19/15 10:33:47	admin	UPDATED	BwEngineCpuUsedHigh	Default	Default	5
09/19/15 10:33:36	admin	UPDATED	BwEngineCpuUsedHigh	Default	Default	0.0
09/19/15 10:33:21	admin	UPDATED	BwActivityExecutionTimeHig	Default	Default	5
09/19/15 10:33:06	admin	UPDATED	BwActivityExecutionTimeHig	Default	Default	0.0
09/19/15 10:32:50	admin	UPDATED	BwActivityErrorRateHigh	Default	Default	5
09/19/15 10:32:19	admin	UPDATED	BwActivityErrorRateHigh	Default	Default	0.0
09/19/15 09:42:07	admin	UPDATED	BwEngineCpuUsedHigh	Default	Default	5
09/19/15 09:41:42	admin	UPDATED	BwActivityExecutionTimeHig	Default	Default	5
09/19/15 09:41:30	admin	UPDATED	BwActivityExecutionTimeHig	Default	Default	0.0
09/19/15 09:40:59	admin	UPDATED	BwActivityErrorRateHigh	Default	Default	5
09/19/15 09:40:30	admin	UPDATED	BwActivityErrorRateHigh	Default	Default	0.0
09/19/15 09:39:30	admin	UPDATED	BwActivityExecutionTimeHig	Default	Default	5
09/19/15 09:39:09	admin	UPDATED	BwActivityExecutionTimeHig	Default	Default	5
09/19/15 09:34:23	admin	UPDATED	BwActivityExecutionTimeHig	Default	Default	0.0
09/19/15 09:34:07	admin	UPDATED	BwActivityErrorRateHigh	Default	Default	5

Title Bar (possible features are):

Open the previous and upper display.

Open an instance of this display in a new window.

Open the online help page for this display.

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:

Audit Conn OK	The Alert Server connection state. Disconnected. Connected.
TIME_STAMP	The date and time of the modification.
USER	The user name of the administrator who made the modification.
ACTION	The type of modification made to the alert, such as **UPDATED**.
ALERTNAME	The name of the alert modified.

INDEXTYPE	The type of alert Index. Index Type refers to the manner in which alert settings are applied and vary among CI Types. For example, the JVM CI Type has a PerJvm Index Type, the EMS CI Type has PerServer, PerTopic and PerQueue Index Types which apply alerts to servers, topics and queues, respectively.
ALERTINDEX	The index of the alert which identifies its source.
WARNINGLEVEL	The warning threshold value for the alert at the time this modification was made, as indicated in the **TIME_STAMP** column. The warning level is a threshold that, when exceeded, a warning is executed.
ALARMLEVEL	The alarm threshold value for the alert at the time this modification was made, as indicated in the **TIME_STAMP** column. The alarm level is a threshold that, when exceeded, an alarm is executed.
DURATION	The duration value for the alert at the time this modification was made, as indicated in the **TIME_STAMP** column. The alert duration is the amount of time (in seconds) that a value must be above the specified Warning Level or Alarm Level threshold before an alert is executed. **0** is for immediate execution.
ENABLED	When checked, indicates the alert was enabled at the time this modification was made, as indicated in the **TIME_STAMP** column.
USEINDEX	When checked, indicates the alert override was enabled at the time this modification was made, as indicated in the **TIME_STAMP** column. For details about alert overrides, see **Alert Administration**.

RTView Agent Metrics Administration

Verify when agent metrics were last queried by the Monitor. The data in this display is predominantly used for debugging by Technical Support.

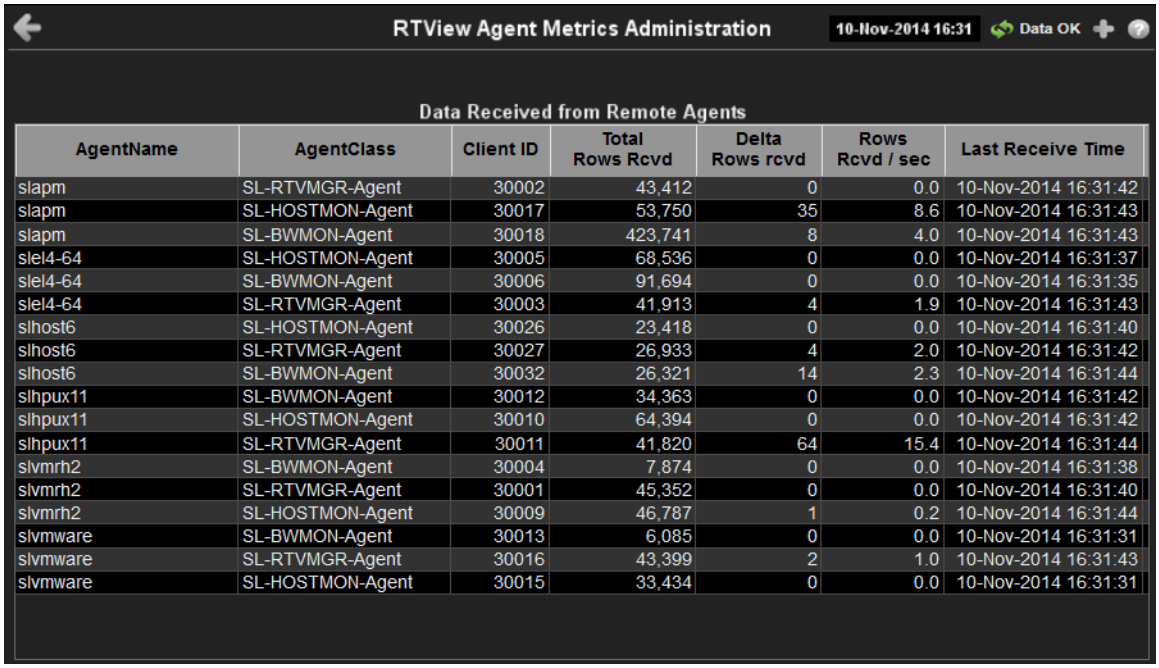

AgentName	AgentClass	Client ID	Total Rows Rcvd	Delta Rows rcvd	Rows Rcvd / sec	Last Receive Time
slapm	SL-RTVMGR-Agent	30002	43,412	0	0.0	10-Nov-2014 16:31:42
slapm	SL-HOSTMON-Agent	30017	53,750	35	8.6	10-Nov-2014 16:31:43
slapm	SL-BWMON-Agent	30018	423,741	8	4.0	10-Nov-2014 16:31:43
slel4-64	SL-HOSTMON-Agent	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

Title Bar (possible features are):

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.

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.

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.

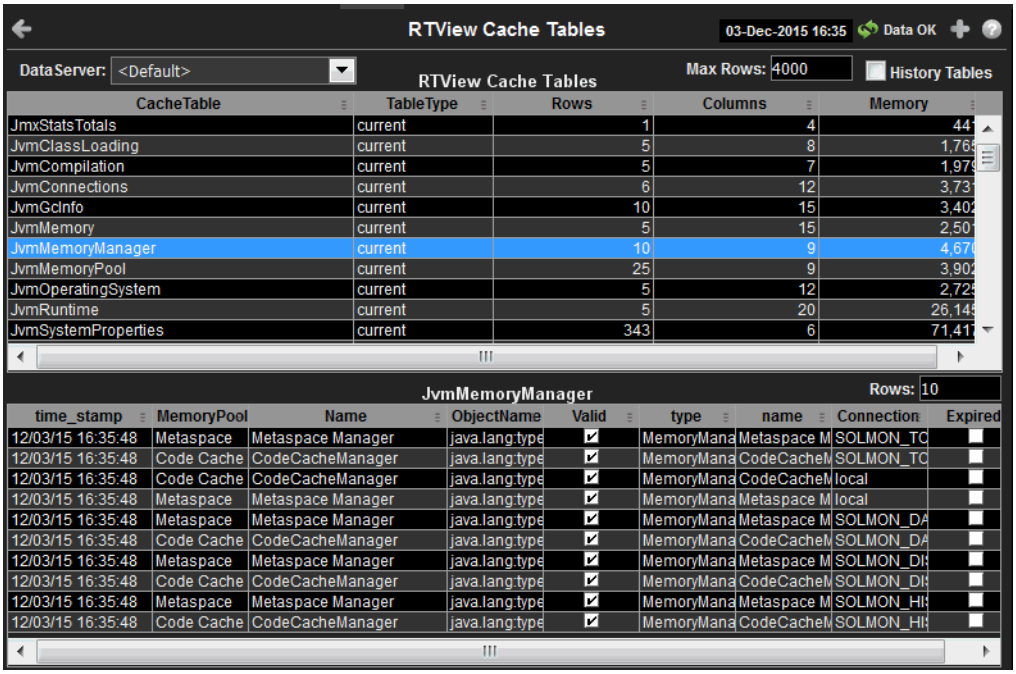

Title Bar (possible features are):

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.	

RTView Agent Administration

Verify when agent metrics were last queried by the Monitor. The data in this display is predominantly used for debugging by Technical Support.

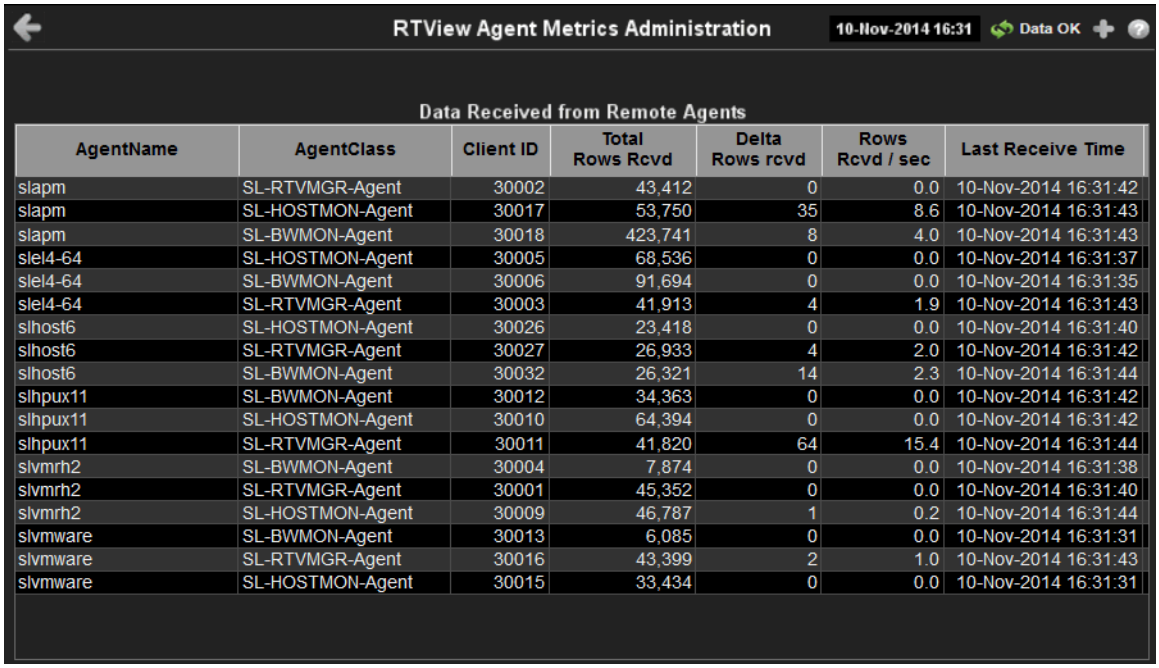

AgentName	AgentClass	Client ID	Total Rows Rcvd	Delta Rows rcvd	Rows Rcvd / sec	Last Receive Time
slapm	SL-RTVMGR-Agent	30002	43,412	0	0.0	10-Nov-2014 16:31:42
slapm	SL-HOSTMON-Agent	30017	53,750	35	8.6	10-Nov-2014 16:31:43
slapm	SL-BWMON-Agent	30018	423,741	8	4.0	10-Nov-2014 16:31:43
slel4-64	SL-HOSTMON-Agent	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

Title Bar (possible features are):

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.

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.

About

Verify when agent metrics were last queried by the Monitor. The data in this display is predominantly used for debugging by Technical Support.

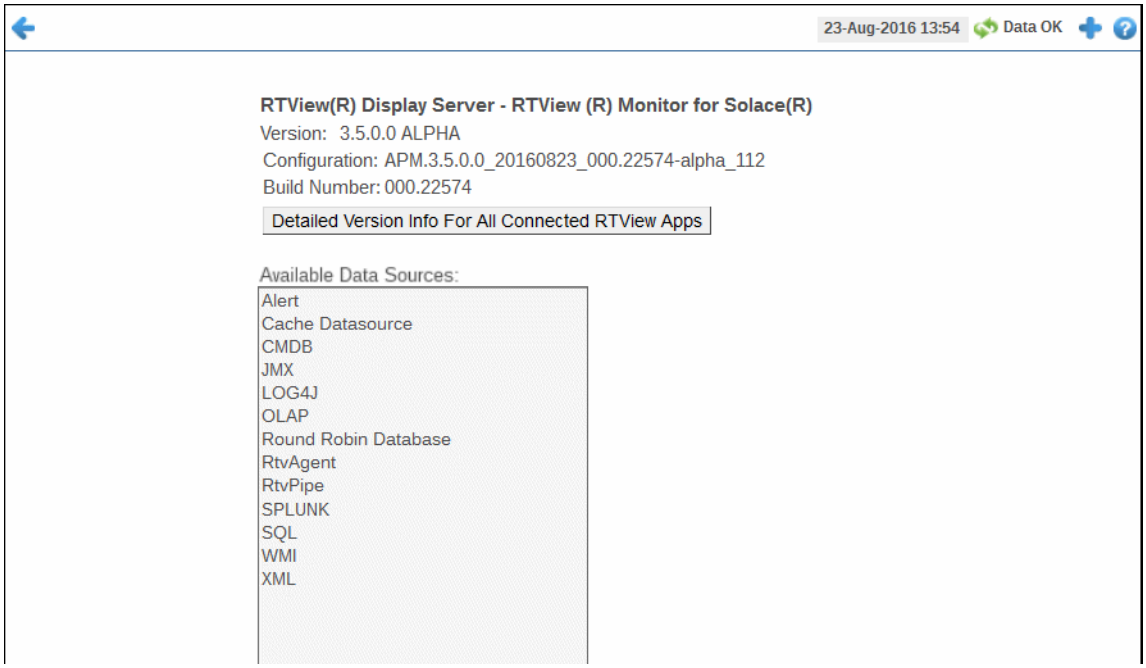

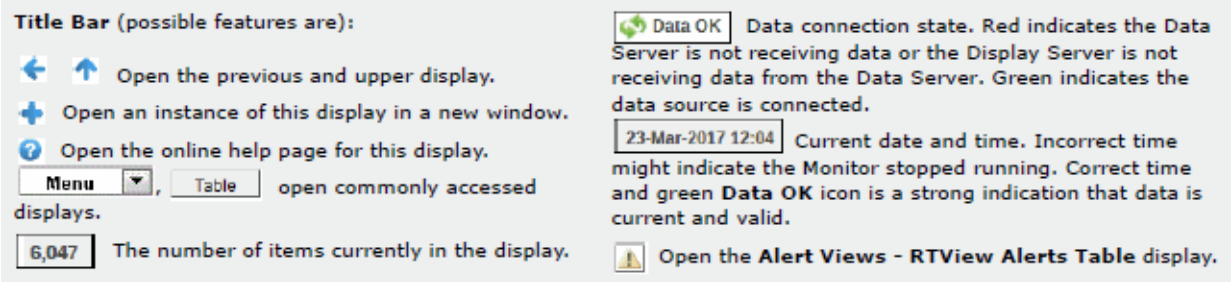

Data Received from Remote Agents Table	
AgentName	Name of the agent.
AgentClass	Class of the agent.
Client ID	Unique client identifier.
Total Rows Rcvd	Total number of rows of data received.
Rows Rcvd/sec	Number of rows of data received per second.
Last Receive Time	Last time data was received from the agent.

RTView Manager Views/Displays

This section describes RTView Manager displays. Use RTView Manager displays to track the health of Solace Monitor components.

Note that the "MySQL Database" and "Docker Engines" displays are populated with performance data only if you are using the RTView Monitor for Solace AMI version.

The RTView Manager has the following Views:

- "JVM Process Views"
- "RTView Servers"
- "Tomcat Servers"
- "MySQL Database"
- "Docker Engines"
- "Hosts"
- "Alert Views"
- "Administration"

JVM Process Views

These displays present performance data for monitored Java Virtual Machine (JVM) Processes. Use these displays to examine the performance and resource use of JVMs in summary and detail. Any JVM that is enabled for monitoring can be included in these displays. The displays include summary overviews and detail pages with historical trends.

You can set alert thresholds on performance and resource metrics for your JVMs, including **CPU Percent**, **Memory Used** and **Gc Duty cycle**. Alerts are shown in the "All JVMs Heatmap" display. Use the detailed JVM displays to investigate further; for example a **Memory Used** alarm might take you to the "JVM Summary" display to get historical memory use, or a **Gc Duty Cycle** alarm might take you to the "JVM GC Trends" display.

Displays in this View are:

- "All JVMs Heatmap": Heatmap of alert states for all JVM connections

- "All JVMs Table": Table of connection details for all JVM connections.
- "JVM Summary": Table of connection details for a single JVM as well as performance trend graphs.
- "JVM System Properties": Table of system details for a single JVM.
- "JVM Memory Pool Trends": Trend graphs of memory pool utilization.
- "JVM GC Trends": Trend graphs of garbage collection memory utilization.

All JVMs Heatmap

View the most critical alert state for all monitored JVM connections for one or all sources, as well as CPU and memory utilization. The heatmap organizes JVM connections by source and host, and uses color to show the most critical Metric value for each JVM connection associated with the selected source. Each rectangle in the heatmap represents a JVM connection. The rectangle size represents the amount of memory reserved for that process; a larger size is a larger value. Each Metric (selected from the drop-down menu) has a color gradient bar that maps relative values to colors.

Choose one or **All Sources** from the **Source** drop-down menu. Use the check-boxes ☑ to include or exclude labels in the heatmap. Move your mouse over a rectangle to see detailed JVM connection information (including **PID**). Drill-down and investigate by clicking a rectangle in the heatmap to view details for the selected connection in the **JVM Summary** display.

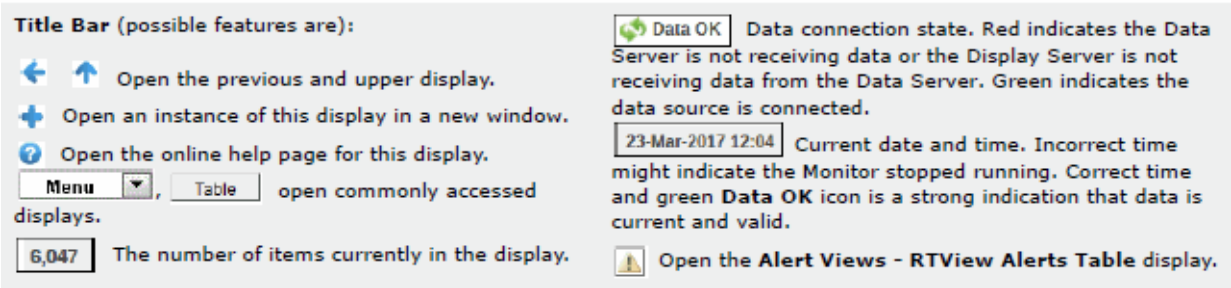

Fields and Data
This display includes:

Source Choose one or **All Sources** to display.

JVM Count The number of JVM connections shown in the display.

Show Inactive Select to show inactive connections.

Connection Select to show JVM connections names.

Metric
Select the Metric to display in the heatmap. Each Metric has a color gradient bar that maps relative values to colors.

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

Red indicates that one or more alerts have reached their alarm threshold. Alerts that have exceeded their specified ALARM LEVEL threshold have an Alert Severity value of **2**.

Yellow indicates that one or more alerts have reached their alarm threshold. Alerts that have exceeded their specified WARNING LEVEL threshold have an Alert Severity value of **1**.

Green indicates that no alerts have reached their alert thresholds. Alerts that have not exceeded their specified thresholds have an Alert Severity value of **0**.

Alert Count The number of alerts for the rectangle. The color gradient bar values range from **0** to the maximum number of alerts in the heatmap.

CPU % The total percent (%) CPU utilization for the rectangle. The color gradient bar values range from **0** to the maximum percent (%) CPU utilization in the heatmap.

Memory % The total percent (%) memory utilization for the rectangle. The color gradient bar values range from **0** to the maximum percent (%) memory utilization in the heatmap.

Current Heap The current amount of heap committed for the connection, in kilobytes. The color gradient bar values range from **0** to the maximum amount in the heatmap.

Used Heap The total amount of heap used by the connection, in kilobytes. The color gradient bar values range from **0** to the maximum amount used in the heatmap.

All JVMs Table

View JVM connection details for one or all sources, the most critical alert state for each JVM connection, as well as CPU and memory utilization in a tabular format. Each row in the table is a different connection.

Choose one or **All Sources** from the **Source** drop-down menu. Check the **Show Inactive** box to include inactive connections. The row color for inactive connections is dark red. Click Sort to order column data. Drill-down and investigate by clicking a row in the table to view details for the selected connection in the **JVM Summary** display.

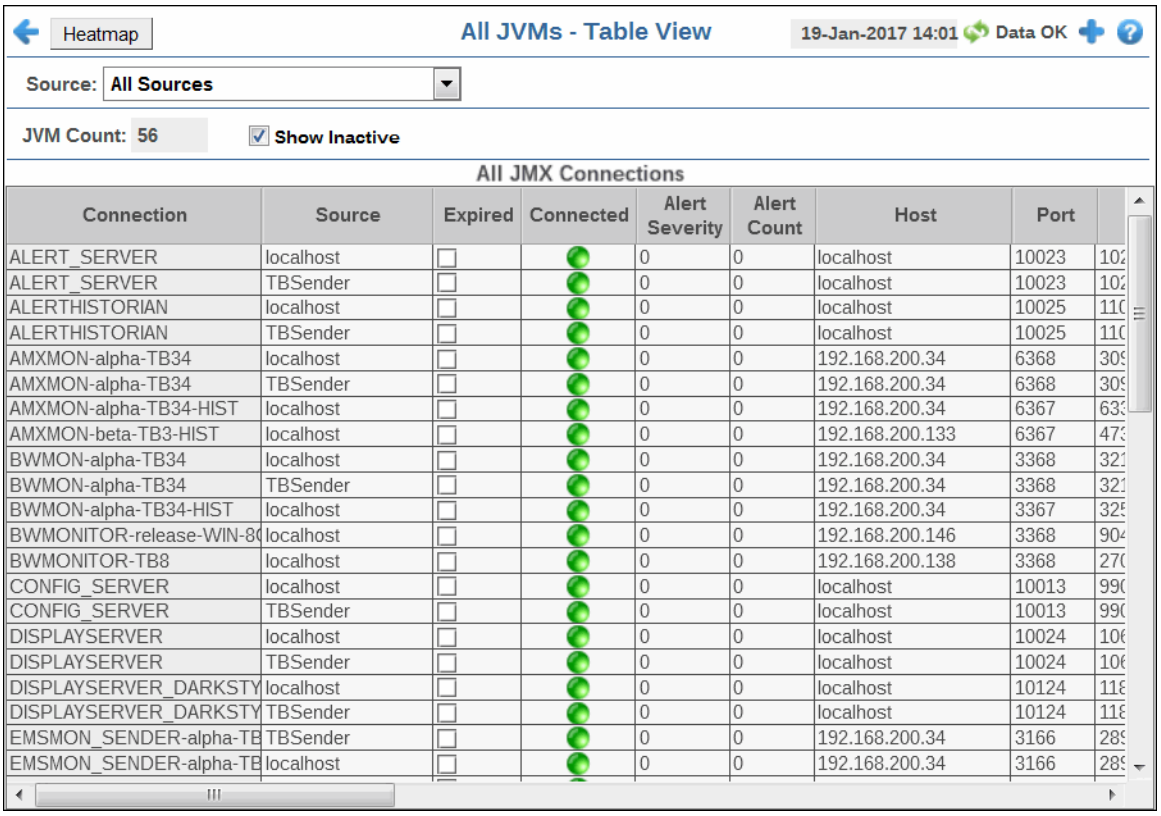

Title Bar (possible features are):

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:

Source Choose one or **All Sources** to display.

JVM Count: The number of JVM connections in the table.

Show Inactive Select to include inactive connections.

All JMX Connections Table

Connection	The name of the JVM connection.
Source	The name of the source.
Expired	When checked, this connection is expired due to inactivity.
Connected	The data connection state: Disconnected. Connected.
Alert Severity	The maximum level of alerts associated with the connection. Values range from **0** to **2**, where **2** is the greatest Alert Severity. One or more alerts associated with the connection exceeded their ALARM LEVEL threshold. One or more alerts associated with the connection exceeded their WARNING LEVEL threshold. No alerts associated with the connection have exceeded their thresholds.
Alert Count	The current number of alerts for the connection.
Host	The name of the host for this connection.
Port	The port number for the connection.
PID	The connection PID.
CPU %	The amount of CPU, in percent (%) used by this connection.
Max Heap	The maximum amount of heap used by this connection, in kilobytes.
Current Heap	The current amount of committed heap for this connection, in kilobytes.
Used Heap	The current amount of heap used by this connection, in kilobytes.
Mem % Used	The amount of JVM memory used by this connection, in percent (%).
RtvAppType	The type of RTView application, where: **1** is for the Historian, **3** is for the Data Server; **5** is for the Display Server, and **0** is a non-RTView application.
Source	The Data Server that sent this value.
time_stamp	The date and time this row of data was last updated.

JVM Summary

Track JVM memory and CPU usage, get JVM system information, application performance metrics, and input arguments for a single connection. Verify whether the memory usage has reached a plateau. Or, if usage is getting close to the limit, determine whether to allocate more memory.

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

Title Bar (possible features are):

Open the previous and upper display.

Open an instance of this display in a new window.

Open the online help page for this display.

Menu, Table open commonly accessed displays.

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:

Source Select the type of connection to the RTView Server.

Connection Select an RTView Server from the drop-down menu. Names can be modified in the RTView Server configuration properties file.

Operating System
Displays data pertaining to the operating system running on the host on which the JVM resides.

	Connected	The data connection state: Disconnected. Connected.
	Expired	When checked, this server is expired due to inactivity.
	Operating System	The name of the operating system running on the host on which the JVM resides.
	OS Version	The operating system version.
	Architecture	The ISA used by the processor.
	Available Processors	The total number of processors available to the JVM.
Runtime		
	Process Name	Name of the process.
	Start Time	The date and time that the application started running.
	Up Time	The amount of time the application has been running, in the following format: **0d 00:00** **<days>d <hours>:<minutes>:<seconds>** For example: **10d 08:41:38**
	JVM CPU %	The amount of CPU usage by the JVM, in percent.
	Live Threads	The total number of live threads.
	Daemon Threads	The total number of live daemon threads.
	Peak Threads	The total number of peak live threads since the JVM started or the peak was reset.
	Max Heap Mb	The maximum amount of memory used for memory management by the application in the time range specified. This value may change or be undefined. NOTE: A memory allocation can fail if the JVM attempts to set the **Used** memory allocation to a value greater than the **Committed** memory allocation, even if the amount for **Used** memory is less than or equal to the *Maximum* memory allocation (for example, when the system is low on virtual memory).
	Committed Mb	The amount of memory, in megabytes, guaranteed to be available for use by the JVM. The amount of committed memory can be a fixed or variable size. If set to be a variable size, the amount of committed memory can change over time, as the JVM may release memory to the system. This means that the amount allocated for **Committed** memory could be less than the amount initially allocated. **Committed** memory will always be greater than or equal to the amount allocated for **Used** memory.
	Used Mb	The amount of memory currently used by the application. Memory used includes the memory occupied by all objects including both reachable and unreachable objects.
Class Name	Class name used for JVM.	
Arguments	The arguments used to start the application.	

More Arguments Additional arguments used to start the application.

JVM CPU, Memory, Thread Trends
Shows JVM metrics for the selected server.

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

Base at Zero Use zero as the Y axis minimum for all graph traces.

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

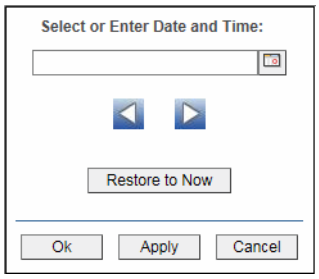

By default, the time range end point is the current time. To change the time range end point, click Calendar and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

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

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

JVM CPU % Traces the amount of memory, in percent, used by the JVM in the time range specified.

Max Heap Mb Traces the maximum amount of memory used for memory management by the application in the time range specified. This value may change or be undefined.

NOTE: A memory allocation can fail if the JVM attempts to set the **Used** memory allocation to a value greater than the **Committed** memory allocation, even if the amount for **Used** memory is less than or equal to the **Maximum** memory allocation (for example, when the system is low on virtual memory).

Cur Heap Mb Traces the current amount of memory, in megabytes, used for memory management by the application in the time range specified.

Used Heap Mb Traces the memory currently used by the application.

Live Threads Traces the total number of currently active threads in the time range specified.

JVM System Properties

Track JVM input arguments and system properties for a single connection. Use the available drop-down menus or right-click to filter data shown in the display.

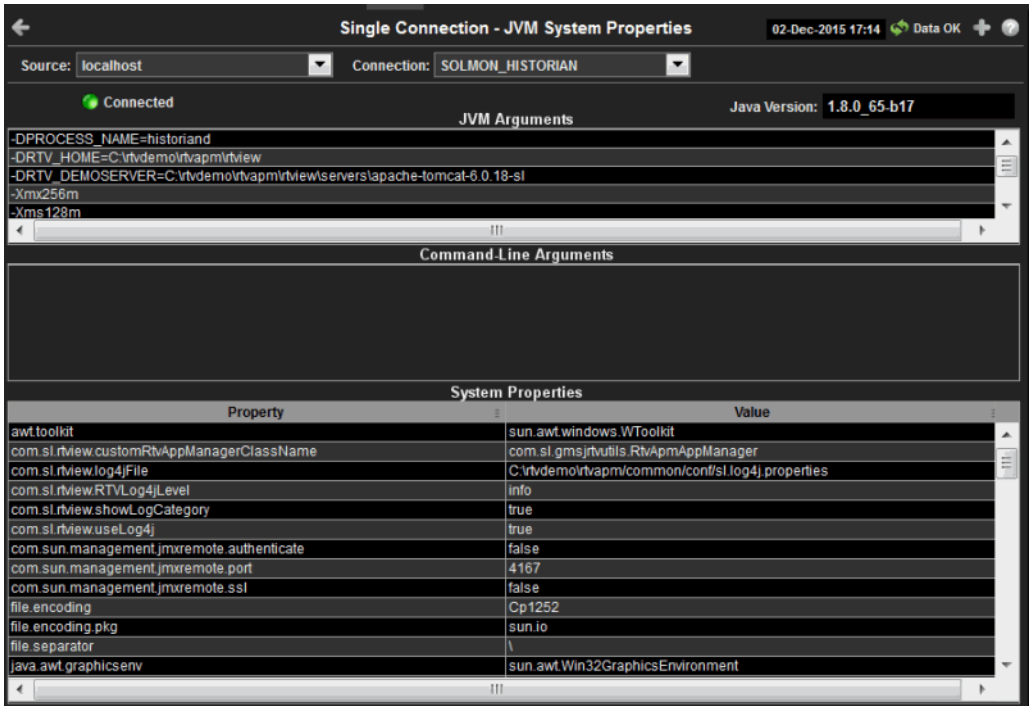

Title Bar (possible features are):

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:

Source	Select the type of connection to the RTView Server.
Connection	Select an RTView Server from the drop-down menu. Names can be modified in the RTView Server configuration properties file.
Connected	The data connection state: Disconnected. Connected.
Java Version	The Java version running on the selected server.
JVM Arguments	The JVM arguments in the **RuntimeMXBean InputArguments** attribute.

Command Line Arguments Arguments used to start the application.

System Properties
This table lists and describes system property settings.

Property Name of the property.

Value Current value of the property.

JVM Memory Pool Trends

Track JVM heap and non-heap memory usage for a single connection. Use the available drop-down menus or right-click to filter data shown in the display.

Title Bar (possible features are):

Open the previous and upper display.

Open an instance of this display in a new window.

Open the online help page for this display.

Menu, Table open commonly accessed displays.

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:

Source Select the type of connection to the RTView Server.

Connection Select an RTView Server from the drop-down menu. Names can be modified in the RTView Server configuration properties file.

Connected The data connection state:
- Disconnected.
- Connected.

Base at Zero Use zero as the Y axis minimum for all graph traces.

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

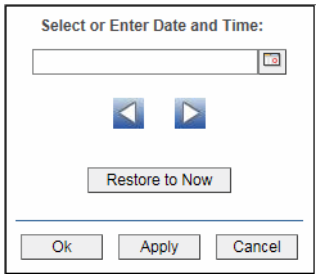

By default, the time range end point is the current time. To change the time range end point, click Calendar and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

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

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

Heap Memory

Maximum	The maximum amount of memory used, in megabytes, for memory management by the application in the time range specified. This value may change or be undefined. NOTE: A memory allocation can fail if the JVM attempts to set the **Used** memory allocation to a value greater than the **Committed** memory allocation, even if the amount for **Used** memory is less than or equal to the **Maximum** memory allocation (for example, when the system is low on virtual memory).
Committed	The amount of memory, in megabytes, guaranteed to be available for use by the JVM. The amount of committed memory can be a fixed or variable size. If set to be a variable size, the amount of committed memory can change over time, as the JVM may release memory to the system. This means that the amount allocated for **Committed** memory could be less than the amount initially allocated. **Committed** memory will always be greater than or equal to the amount allocated for **Used** memory.
Used	The amount of memory, in megabytes, currently used by the application. Memory used includes the memory occupied by all objects including both reachable and unreachable objects.
Peak Tenured Used	The amount of memory, in megabytes, used by tenured JVM objects in the time range specified. Tenured refers to JVM objects contained in a pool that holds objects that have avoided garbage collection and reside in the survivor space. Peak tenured refers to the maximum value of the tenured memory over a specified period of time.
Eden Space	Traces the amount of memory used by the JVM eden pool in the time range specified. Eden refers to the JVM eden pool, which is used to initially allocate memory for most objects.
Survivor Space	Traces the amount of memory used by the JVM survivor pool in the time range specified. The JVM survivor pool holds objects that survive the eden space garbage collection.
Tenured Gen	Traces the amount of memory used by tenured JVM objects in the time range specified. Tenured refers to JVM objects contained in a pool that holds objects that have avoided garbage collection and reside in the survivor space. Peak tenured refers to the maximum value of the tenured memory over a specified period of time.

Non-Heap Memory

Maximum	The maximum amount of memory, in megabytes, used for JVM non-heap memory management by the application in the time range specified.
Committed	The amount of memory, in megabytes, guaranteed to be available for use by JVM non-heap memory management. The amount of committed memory can be a fixed or variable size. If set to be a variable size, it can change over time, as the JVM may release memory to the system. This means that the amount allocated for **Committed** memory could **be** less than the amount initially allocated. Committed memory will always be greater than or equal to the amount allocated for **Used** memory.
Used	The amount of memory, in megabytes, currently used by the application. Memory used includes the memory occupied by all objects including both reachable and unreachable objects.
Objects Pending Finalization	The value of the **MemoryMXBean ObjectPendingFinalizationCount** attribute.
Verbose	The value of the **MemoryMXBean Verbose** attribute.
Code Cache	Traces the amount of non-heap memory used in the JVM for compilation and storage of native code.
Perm Gen	Traces the amount of memory used by the pool containing reflective data of the virtual machine, such as class and method objects. With JVMs that use class data sharing, this generation is divided into read-only and read-write areas.

Operations

Run Garbage Collector	Performs garbage collection on the selected server.
Reset Peak Usage	Clears peak usage on the selected server.

JVM GC Trends

Track JVM garbage collection memory usage for a single connection. Use the available drop-down menus or right-click to filter data shown in the display.

Title Bar (possible features are):

Open the previous and upper display.

Open an instance of this display in a new window.

Open the online help page for this display.

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:

Source	Select the type of connection to the RTView Server.
Connection	Select an RTView Server from the drop-down menu. Names can be modified in the RTView Server configuration properties file.
Garbage Collector	Select a garbage collection method: **Copy** or **MarkSweepCompact**.
Max	Shows the maximum amount of memory used for JVM garbage collection in the time range specified.
Committed	Shows the amount of memory guaranteed to be available for use by JVM non-heap memory management. The amount of committed memory can be a fixed or variable size. If set to be a variable size, it can change over time, as the JVM may release memory to the system. This means that the amount allocated for **Committed** memory could be less than the amount initially allocated. **Committed** memory will always be greater than or equal to the amount allocated for **Used** memory.

Base at Zero Use zero as the Y axis minimum for all graph traces.

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

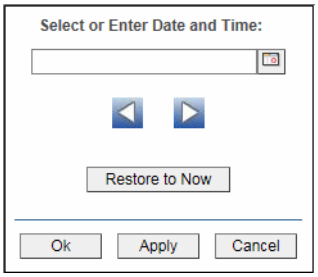

By default, the time range end point is the current time. To change the time range end point, click Calendar and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

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

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

Memory Usage (in MB) Before and After Garbage Collection

Maximum Traces the maximum amount of memory used by garbage collection in the time range specified. This value may change or be undefined.

NOTE: A memory allocation can fail if the JVM attempts to set the **Used** memory allocation to a value greater than the **Committed** memory allocation, even if the amount for **Used** memory is less than or equal to the **Maximum** memory allocation (for example, when the system is low on virtual memory).

Committed Traces the amount of memory guaranteed to be available for use by the JVM. The amount of committed memory can be a fixed or variable size. If set to be a variable size, the amount of committed memory can change over time, as the JVM may release memory to the system. This means that the amount allocated for **Committed** memory could be less than the amount initially allocated. **Committed** memory will always be greater than or equal to the amount allocated for **Used** memory.

Used - Before Traces the amount of memory used before the last garbage collection.

Used - After Traces the amount of memory used after the last garbage collection.

Duration The duration, in seconds, of garbage collection.

Duty Cycle The percentage of time that the application spends in garbage collection.

RTView Servers

These displays present performance data for all RTView Servers. Displays in this View are:

- "Data Servers": Shows metrics for RTView Data Servers.
- "Display Servers": Shows metrics for RTView Display Servers.
- "Historian Servers": Shows metrics for RTView Historian Servers.
- "Version Info": Shows the version information.

Data Servers

Track data transfer metrics for RTView Data Servers, client count and throughput trends. Use the available drop-down menus or right-click to filter data shown in the display.

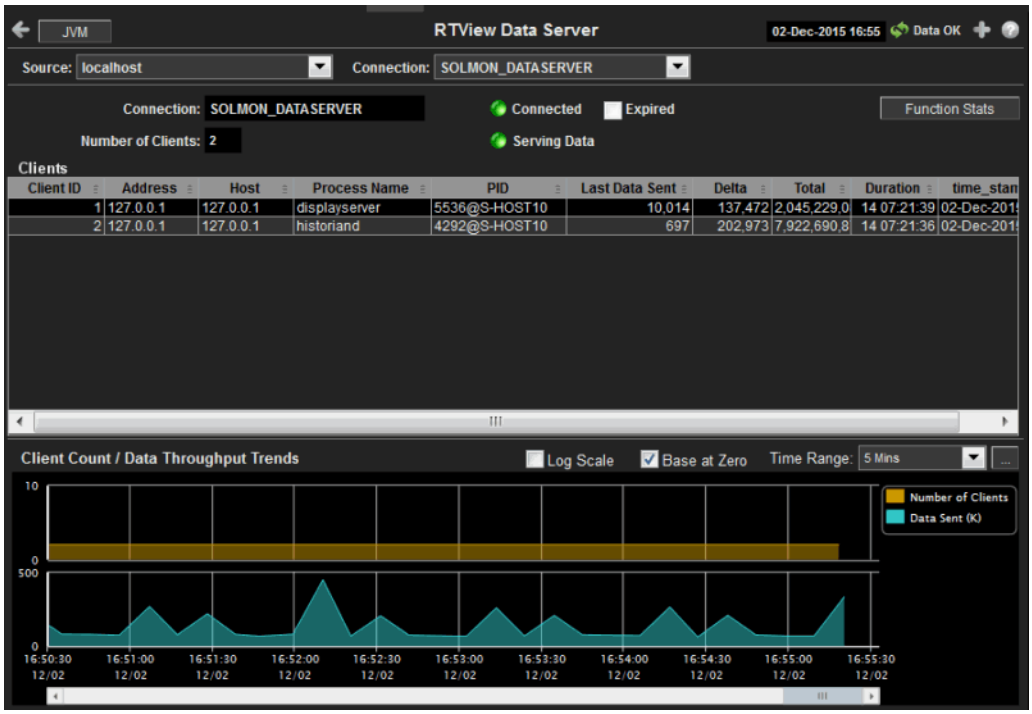

Title Bar (possible features are):

Open the previous and upper display.

Open an instance of this display in a new window.

Open the online help page for this display.

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.

Source	Select the type of connection to the RTView Server.
Connection	Select an RTView Server from the drop-down menu. Names can be modified in the RTView Server configuration properties file.

Connection	The connection selected from the **Connection** drop-down menu.
Number of Clients	The number of clients currently server on this Data Server.
Connected	The Data Server connection state: Disconnected. Connected.
Serving Data	The Data Server is not currently serving data. The Data Server is currently serving data.
Expired	This server has been marked as expired after no activity.
Function Stats	Opens the **RTView Function Stats** display which shows detailed performance statistics for RTView functions in the selected Data Server. This button is only enabled if the RTVMGR has a JMX connection defined for the selected Data Server.

Clients
This table describes all clients on the selected server.

Address	The client IP address.
Client ID	The unique client identifier.
Duration	The amount of time for this client session. Format: **dd HH:MM:SS** **<days> <hours>:<minutes>:<seconds>** **For example:** **10d 08:41:38**
Host	The client host name.
Last Data Sent	The amount of data, in bytes, last sent to the client.
Delta	The amount of data, in bytes, sent since the last update.
Total	The total amount of data, in bytes, sent to the client.
TIME_STAMP	The date and time this row of data was last updated.

Client Count / Data Throughput Trends
Shows throughput metrics for all clients on the selected server.

Log Scale	Enable to use a logarithmic scale for the Y axis. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.
Base at Zero	Use zero as the Y axis minimum for all graph traces.

Time Range

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

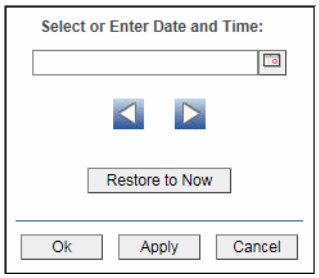

By default, the time range end point is the current time. To change the time range end point, click Calendar and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

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

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

Number of Clients

Traces the number of clients being served by the Data Server.

Data Sent

Traces the total amount of data, in Kilobytes, sent to all clients.

Display Servers

Track display utilization metrics for RTView Display Servers.

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

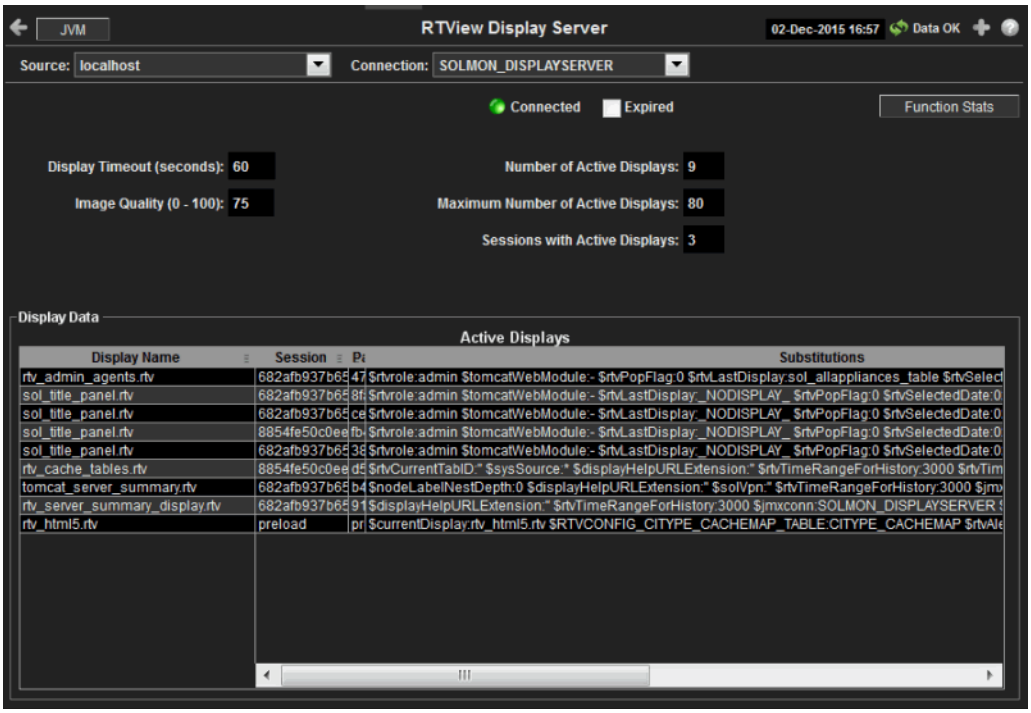

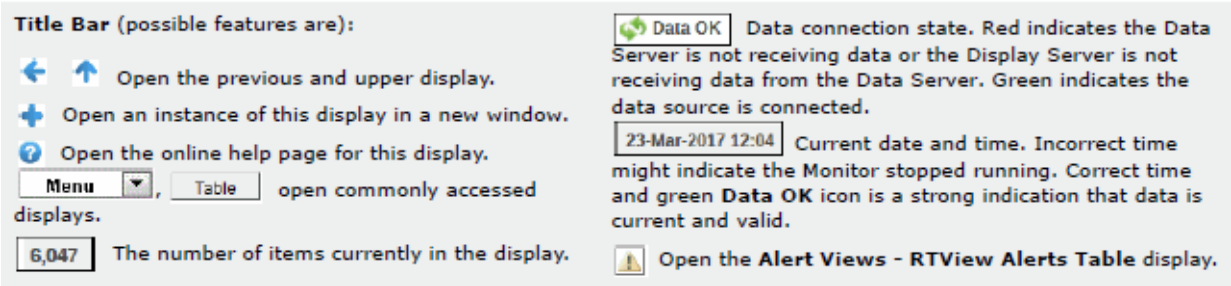

Fields and Data

This display includes:

Source Select the type of connection to the RTView Server.

Connection Select an RTView Server from the drop-down menu. Names can be modified in the RTView Server configuration properties file.

Connected The Display Server connection state:
- Disconnected.
- Connected.

Expired This server has been marked as expired after no activity.

Function Stats	Opens the **RTView Function Stats** display which shows detailed performance statistics for RTView functions in the selected Display Server. This button is only enabled if the RTVMGR has a JMX connection defined for the selected Display Server.
Display Timeout (seconds)	The amount of time, in seconds, that a display can be kept in memory after the Display Servlet has stopped requesting it. The default is **60** seconds (to allow faster load time when switching between displays).
Image Quality (0-100)	A value between **0** and **100**, which controls the quality of the generated images. If the value is **100**, the Display Server outputs the highest quality image with the lowest compression. If the value is **0**, the Display Server outputs the lowest quality image using the highest compression. The default is **75**.
Number of Active Displays	The total number of displays currently being viewed by a user.
Maximum Number of Active Displays	The maximum number of displays kept in memory. The default is **20** (to optimize memory used by the Display Server).
Sessions with Active Displays	Number of clients accessing the Display Server.

Display Data / Active Displays

Display Name	The name of the currently open display.
Session	A unique string identifier assigned to each session.
Panel ID	A unique string identifier assigned to each panel. The Display Server loads each display requested by each client into a panel. This ID can be useful in troubleshooting.
Substitutions	Lists the substitutions used for the display.
Last Ref	The amount of time that has elapsed since the display was last requested by a client.
ID	The client ID.
Preloaded	When checked, indicates that the display (**.rtv**) file is configured in the **DISPLAYSERVER.ini** file to be preloaded. The **history_config** option is used to configure display preloading. Preloading a display makes data immediately available. Preloaded displays are not unloaded unless the Display Server is restarted or the display cache is cleared via JMX. This option can be used multiple times to specify multiple displays to preload.

Historian Servers

Track the status of RTView Historian Servers and data configuration file usage. View the caches that are archived by the Historian application, substitution variables associated with the history cache configuration file, as well as the history cache status. You can also stop and start the Historian, and purge data.

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

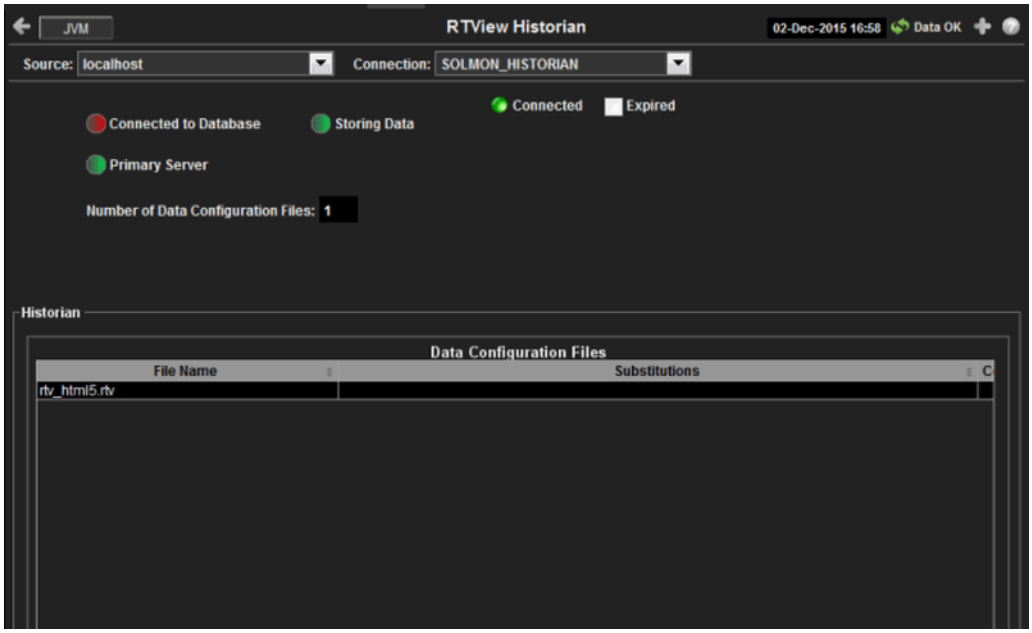

Title Bar (possible features are):

Open the previous and upper display.

Open an instance of this display in a new window.

Open the online help page for this display.

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:

Source	Select the type of connection to the RTView Server.
Connection	Select an RTView Server from the drop-down menu. Names can be modified in the RTView Server configuration properties file.
Connected	The Historian Server connection state: Disconnected. Connected.
Expired	This server has been marked as expired after no activity.
Connected to Database	The Historian Server database connection state: Disconnected. Connected.

Primary Server	When green, indicates that this Historian, when used within a group of Historians, is the primary group member. If the primary member fails or shuts down, the standby member with the highest priority becomes the primary group member. When red, indicates that the Historian is a secondary server. The Historian Server member state: The Historian Server is a secondary group member. This Historian is the primary group member.
Number of Data Configuration Files	The number of configuration files that are used by the history cache.

Historian / Data Configuration Files

File Name	The name of the history cache configuration file.
Substitutions	Lists the substitutions specified in the history cache configuration file.

Version Info

This display provides detailed version information for all of the connected RTView applications. You can view specific applications by filtering data using the **Source**, **Connection**, **Filter Field**, and **Filter Value** fields at the top of the display. This display provides valuable information about the version of each jar that is used in each connected RTView application that can be used to help Technical Support when issues arise. 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.

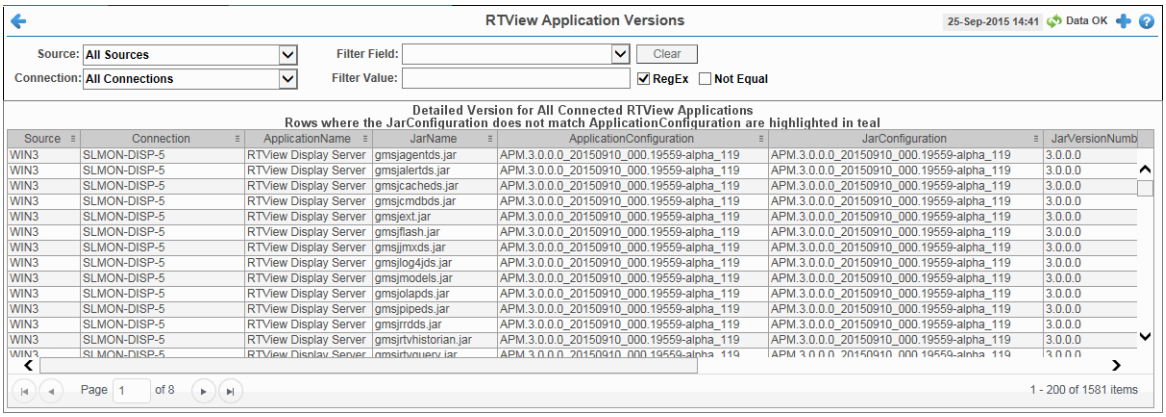

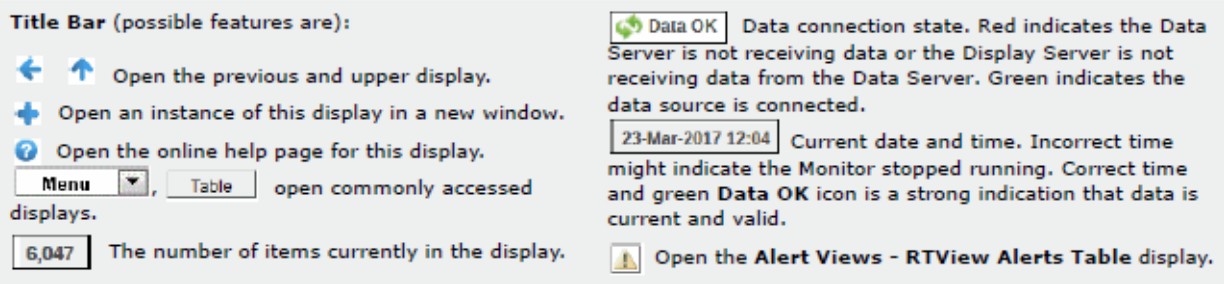

Fields and Data
This display includes:

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.

Tomcat Servers

These displays present performance data for monitored Tomcat Application Servers. Use these displays to examine the state and performance of your Tomcat servers as well as all installed web modules. The server displays include summary overviews and detail pages with historical trends. Displays in this View are:

- "All Tomcat Servers": Table of connection details and performance metrics for all Tomcat connections.
- "Tomcat Server Summary": Performance metrics for one Tomcat Server, including current and historic performance metrics.
- "All Applications Heatmap": Heatmap of performance metrics for all Web modules for one Tomcat Server.
- "Single Application Summary": Table and trend graphs of performance metrics for Web modules.

All Tomcat Servers

View Tomcat Server details per connection such as the total number of sessions, bytes sent/received, and processing time. Each row in the table is a different Tomcat Server. The row color for inactive connections is dark red.

Use this display to get Tomcat server session counts, access and request rates, cache hit rates and data transmission metrics.

Drill-down and investigate by clicking a row in the table to view details for the selected connection in the **Service Summary** display.

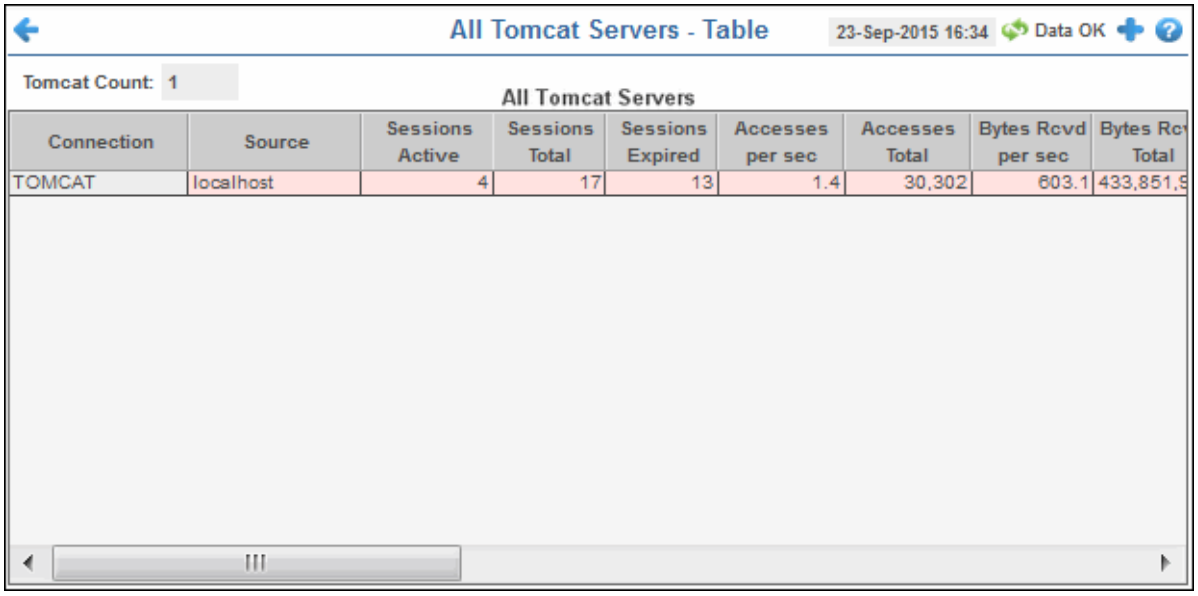

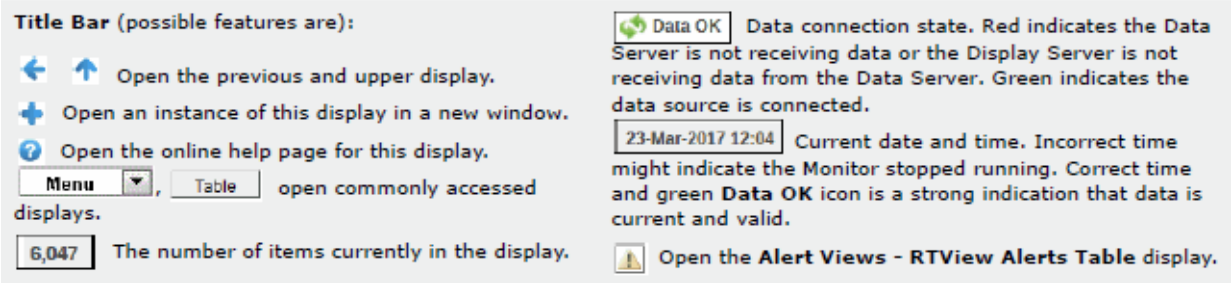

Fields and Data
This display includes:

Tomcat Count	The number of Tomcat connections in the table.
Connection	The name of the Tomcat connection.
Source	The host where the Tomcat Server is running.
Sessions Active	The number of currently active client sessions.
Sessions Total	The total number of client sessions since the server was started.
Sessions Expired	The total number of client sessions that expired since the server was started.
Accesses per sec	The number of times pages are accessed, per second.

Accesses Total	The total number of times pages have been accessed since the server was started.
Bytes Rcvd per sec	The number of bytes received per second.
Bytes Rcvd Total	The total number of bytes received since the server was started.
Bytes Sent per sec	The number of bytes sent per second.
Bytes Sent Total	The total number of bytes sent since the server was started.
Cache Hit Rate	The number of times the cache is accessed, per second.
Requests per sec	The number of requests received, per second.
Requests Total	The total number of requests received since the server was started.
Process Time	The average amount of time, in milliseconds, to process requests.
Error Count	The number of errors that have occurred since the server was started.
appBase	The directory in which Tomcat is installed.
Display Name	The name of the currently open display.
Expired	When checked, this connection is expired due to inactivity.
time_stamp	The date and time this row of data was last updated. Format: **MM/DD/YY HH:MM:SS** **<month>/ <day>/<year> <hours>:<minutes>:<seconds>**

Tomcat Server Summary

Track the performance of one Tomcat Server and get Tomcat hosting and connection details. You can drill down to this display from the Servers table for detailed information and historical trends for a specific server. The trends include Active Sessions, Requests per Sec, and Process Time.

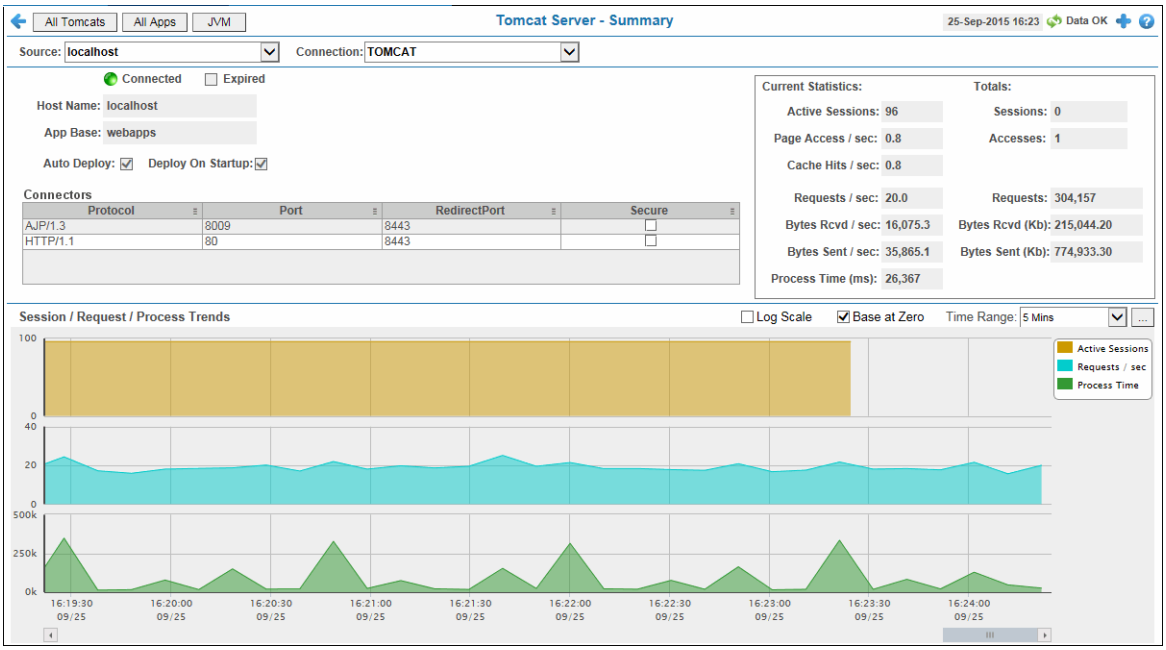

Title Bar (possible features are):

- 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:

Source	Select the host where the Tomcat Server is running.
Connection	Select a Tomcat Server from the drop-down menu.
Connected	The Tomcat Server connection state: Disconnected. Connected.
Expired	When checked, this server is expired due to inactivity.
Host Name	The name of the host where the application resides.
App Base	The directory in which Tomcat modules are installed.

Auto Deploy	When checked, indicates that the Tomcat option, automatic application deployment, is enabled. **Note:** This Tomcat option is set using the **autoDeploy** property in the **server.xml** file, located in the Tomcat **conf** directory. **autoDeploy=true** enables the option.
Deploy On Startup	When checked, indicates that the option to deploy the application on Tomcat startup is enabled. **Note:** This Tomcat option is set using the **deployOnStartup** property in the **server.xml** file, located in the Tomcat **conf** directory. When enabled (**deployOnStartup=true**), applications from the host are automatically deployed.

Connectors
This table shows Tomcat application connection information.

Protocol	The protocol used by the Tomcat application on the host.
Port	The port number used by the Tomcat application on the host.
RedirectPort	The redirect port number used by the Tomcat application on the host.
Secure	When checked, specifies that the Tomcat application uses a secure connection on the host.

Current Statistics / Totals

Active Sessions	The number of clients currently in session with the servlet.
Sessions	The total number of client sessions since the server was started.
Page Access / sec	The number of times pages are accessed, per second.
Accesses	The total number of page accesses since the server was started.
Cache Hits / sec	The number of times the cache is accessed, per second.
Requests / sec	The number of requests received, per second.
Requests	The total number of requests since the server was started.
Bytes Rcvd / sec	The number of bytes received, per second.
Bytes Rcvd (Kb)	The number of kilobytes received since the server was started.
Bytes Sent / sec	The number of bytes sent, per second.
Bytes Sent (Kb)	The total number of kilobytes sent since the server was started.
Process Time	The amount of time, in milliseconds, for the servlet to process client requests.

Session / Request / Process Trends
Shows metrics for the selected server.

Log Scale	Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.
Base at Zero	Use zero as the Y axis minimum for all graph traces.
Time Range	Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar . 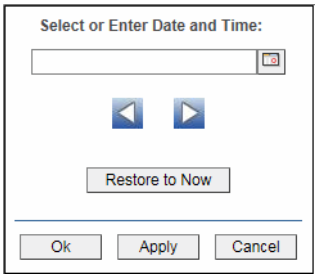 By default, the time range end point is the current time. To change the time range end point, click Calendar and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**. Use the navigation arrows to move forward or backward one time period. NOTE: The time period is determined by your selection from the **Time Range** drop-down menu. Click **Restore to Now** to reset the time range end point to the current time.
Active Sessions	Traces the number of currently active client sessions.
Requests /sec	Traces the number of requests received, per second.
Process Time	Traces the average amount of time, in milliseconds, to process requests.

All Applications Heatmap

View performance metrics for all monitored Tomcat Web modules for one Tomcat Server. The heatmap organizes Tomcat Web modules by server, and uses color to show the most critical Metric value for each Tomcat connection associated with the selected source. Each rectangle in the heatmap represents a Web module. In this heatmap, the rectangle size is the same for all Web modules. Each Metric (selected from the drop-down menu) has a color gradient bar that maps relative values to colors.

Use this display to see at-a-glance the health of all your web applications. You can select the heatmap color metric from a list including active sessions, access rate, and total access count.

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

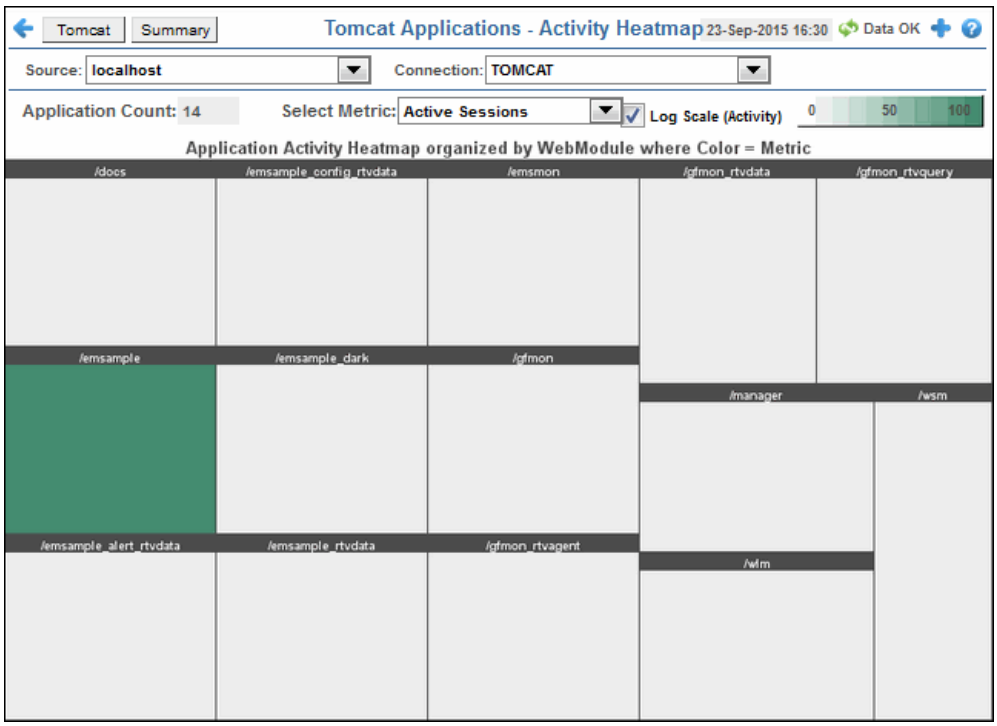

Title Bar (possible features are):

Open the previous and upper display.

Open an instance of this display in a new window.

Open the online help page for this display.

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:

Source	Select the host where the Tomcat Server is running.
Connection	Select a Tomcat Server from the drop-down menu.
Application Count	The number of Tomcat applications in the heatmap.

Log Scale (Activity)	Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.
Select Metric	Select the metric to display in the heatmap. Each Metric has a color gradient bar that maps relative values to colors.

Single Application Summary

Track the performance of all web application modules in a server and view utilization details. The table summarizes the sessions, accesses, cache hit and so forth, for all installed web modules. Each row in the table is a different web application module. The row color for inactive modules is dark red. Select a web application module to view metrics in the trend graph.

Use this data to verify response times of your Web application modules.

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

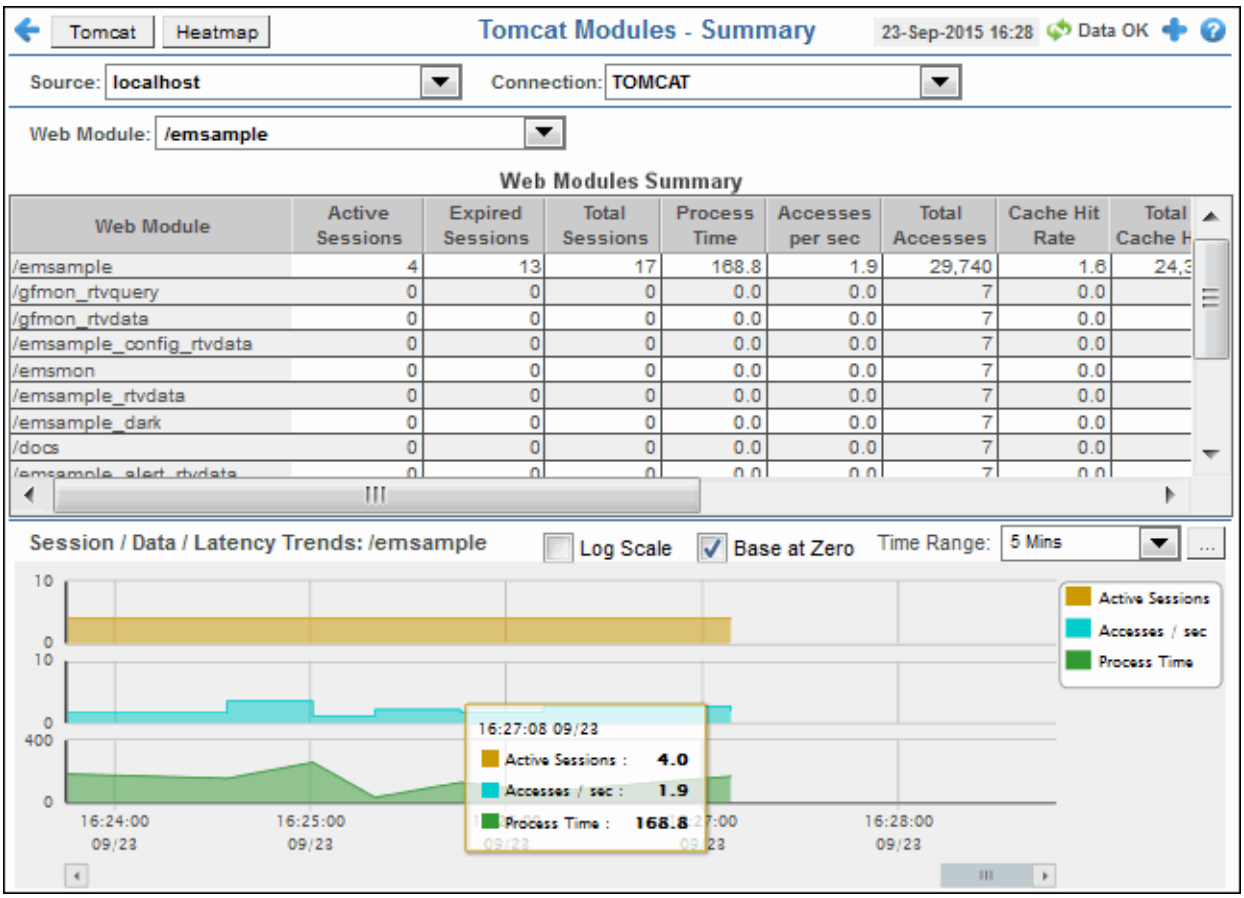

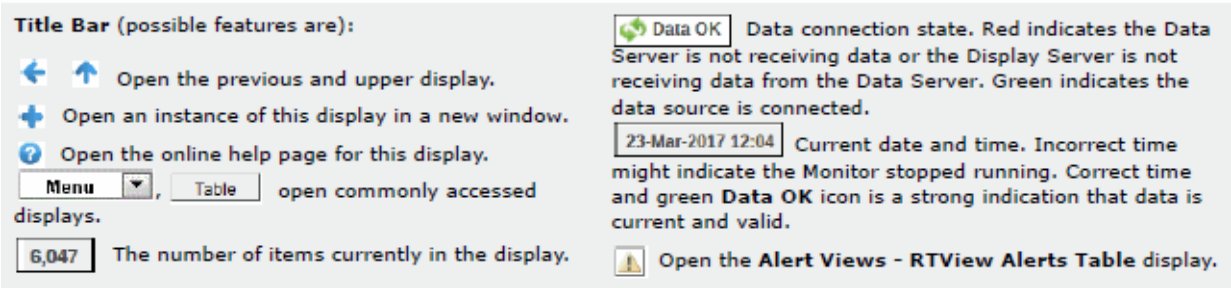

Fields and Data
This display includes:

Source Select the host where the Tomcat Server is running.

Connection Select a Tomcat Server from the drop-down menu. This menu is populated by the selected Source.

Web Module Select a Web module from the drop-down menu. This menu is populated by the selected Connection. The Web Module you select populates the trend graphs.

Web Module Summary

Web Module The name of the Web module.

Sessions Active The number of currently active client sessions.

Sessions Total The total number of client sessions since the application was started.

Sessions Expired The total number of client sessions that expired since the application was started.

Accesses per sec The number of times pages are accessed, per second.

Accesses Total The total number of times pages have been accessed since the application was started.

Bytes Rcvd per sec The number of bytes received per second.

Bytes Rcvd Total The total number of bytes received since the application was started.

Bytes Sent per sec The number of bytes sent per second.

Bytes Sent Total The total number of bytes sent since the application was started.

Cache Hit Rate The number of times the cache is accessed, per second.

Requests per sec The number of requests received, per second.

Requests Total The total number of requests received since the application was started.

Process Time The average amount of time, in milliseconds, to process requests.

Error Count	The number of errors occurred since the application was started.
appBase	The directory in which Tomcat is installed.
Expired	When checked, this connection is expired due to inactivity.
time_stamp	The date and time this row of data was last updated. Format: **MM/DD/YY HH:MM:SS** **<month>/ <day>/<year> <hours>:<minutes>:<seconds>**

Session/Data/Latency Trends
Shows metrics for the selected Web module. The Web module can be selected from the **Web Module** drop-down menu or the **Web Modules Summary** table.

Log Scale	Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.
Base at Zero	Use zero as the Y axis minimum for all graph traces.
Time Range	Select a time range from the drop down menu varying from **2 Minutes** to **Last 7 Days**, or display **All Data**. To specify a time range, click Calendar .

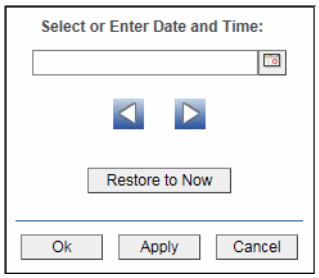

By default, the time range end point is the current time. To change the time range end point, click Calendar and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

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

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

Active Sessions	Traces the number of currently active client sessions.
Accesses / sec	Traces the number of times pages are accessed, per second.
Process Time	Traces the average amount of time, in milliseconds, to process requests.

MySQL Database

The MySQL Database displays provide extensive visibility into the health and performance of the MySQL database included in the RTView Monitor for Solace AMI version. These displays are populated with performance data if you are using the RTView Monitor for Solace AMI version.

Displays in this View are:

- "All Servers Heatmap": A heatmap view of all servers and their associated metrics.
- "All Servers Table": A tabular view of your servers and their associated metrics.
- "Server Summary": Displays performance, processing, alerts, memory, and trend data for a particular database server.
- "Servers Properties": Displays the values of properties on servers.
- "Servers Operations": Trend graph that traces server queries, slow queries, KB sent and KB received.
- "Servers Operations": A tabular view of cache tables performance and utilization metrics.

All Servers Heatmap

This heatmap display provides an easy-to-view interface that allows you to quickly identify the current status of each of your servers. Choose a metric from the **Metric** drop down menu. By default, this display shows the heatmap based on the **Alert Severity** metric. Other metrics are Alert Count, Received, Sent, Delayed Writes, Table Locks Waited, Slow Queries, Slow Launch Threads and Qcache Low Mem Prunes.

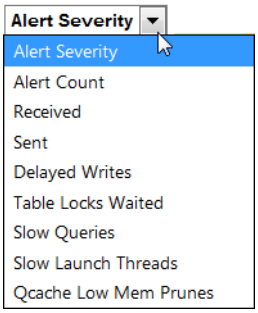

Each rectangle in the heatmap is a different server. Use the **Names** check-box ☑ to include or exclude labels in the heatmap, and mouse over a rectangle to see additional metrics for a server. Click a rectangle to open the "Server Summary" display and see additional details for the selected server.

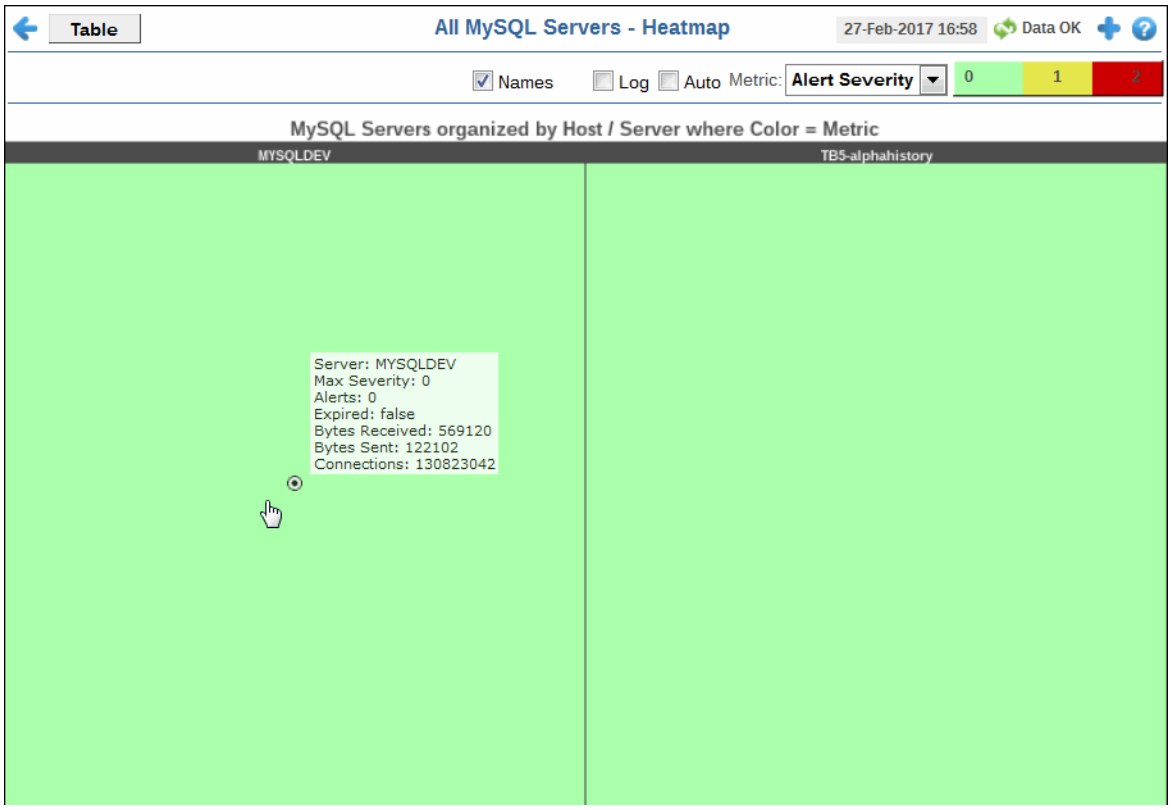

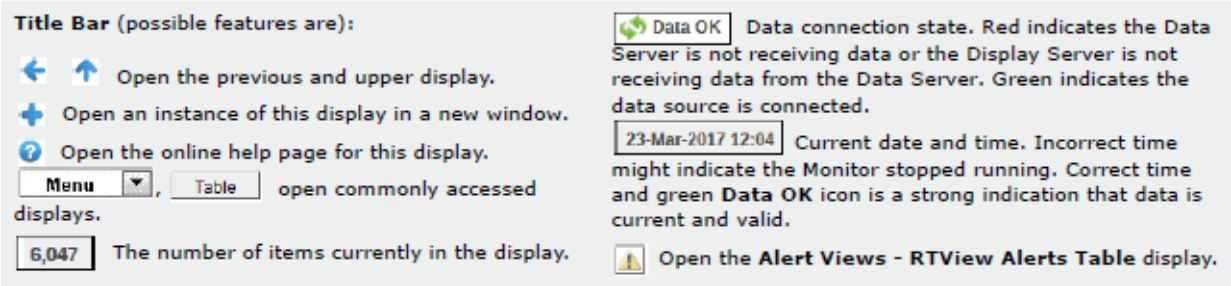

Fields and Data:

Names Select this check box to display the names of the instances at the top of each rectangle in the heatmap.

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

Auto Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value.

Note: Some metrics auto-scale automatically, even when **Auto** is not selected.

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

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

Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.

Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.

Green indicates that no metrics have exceeded their alert thresholds.

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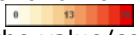 bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.

Received The total number of bytes received. The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the alarm threshold specified for the **MysqlBytesReceivedHigh** alert. The middle value in the gradient bar indicates the average count.

Sent The total number of bytes sent. The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the alarm threshold specified for the **MysqlBytesSentHigh** alert. The middle value in the gradient bar indicates the average count.

Delayed Writes The total number of delayed writes. Values range from **0** to the alarm threshold specified for the **MysqlDelayedWrites** alert. The middle value in the gradient bar indicates the average count:

Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.

Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.

Green indicates that no metrics have exceeded their alert thresholds.

Table Locks Waited The total number of table locks waited. Values range from **0** to the alarm threshold specified for the **MysqlLocksWaited** alert. The middle value in the gradient bar indicates the average count:

Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.

Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.

Green indicates that no metrics have exceeded their alert thresholds.

Slow Queries The total number of slow queries. Values range from **0** to the alarm threshold specified for the **MysqlSlowQueries**. The middle value in the gradient bar indicates the average count:

Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.

Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.

Green indicates that no metrics have exceeded their alert thresholds.

Slow Launch Threads	The total number of slow launch threads. Values range from **0** to the alarm threshold specified for the **MysqlSlowThreads**. The middle value in the gradient bar indicates the average count: Red indicates that one or more metrics exceeded their ALARM LEVEL threshold. Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold. Green indicates that no metrics have exceeded their alert thresholds.
Qcache Low Mem Prunes	The total number of Qcache low memory prunes. Values range from **0** to the alarm threshold specified for the **MysqlQcacheLowMemPrunes**. The middle value in the gradient bar indicates the average count: Red indicates that one or more metrics exceeded their ALARM LEVEL threshold. Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold. Green indicates that no metrics have exceeded their alert thresholds.

All Servers Table

This display provides a tabular view of the performance metrics shown in the "All Servers Heatmap" (alert level, alert count, bytes received, and so forth), as well as additional metrics (such as query information and uptime).

Each table row is a different server. Click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by clicking a row to view details for a server in the "Server Summary" display.

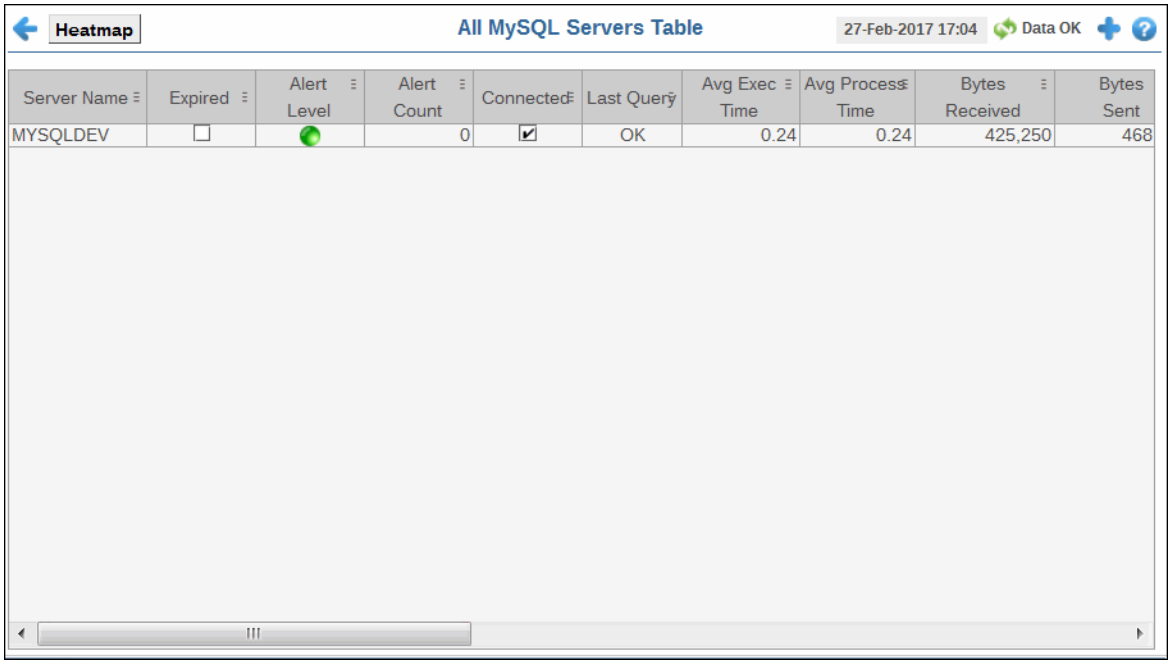

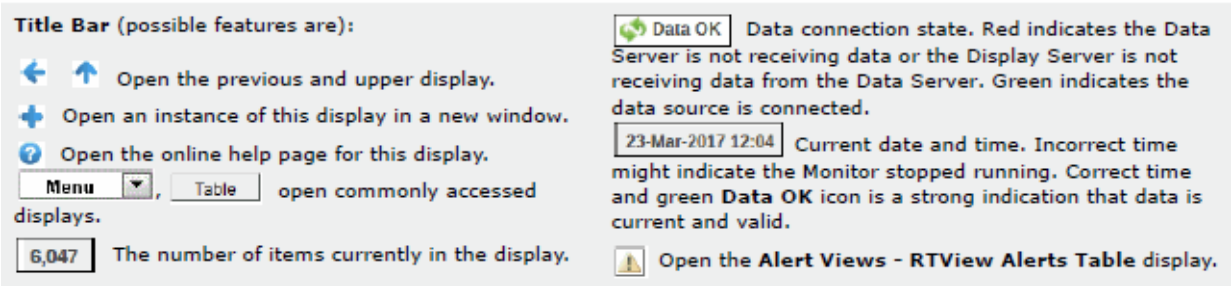

All MySQL Servers Table

Server Name	The name of the server.

Expired	When checked, performance data about the server has not been received within the time specified (in seconds) in the **$mysqlRowExpirationTime** field in the **conf\rtvapm_mysqlmon.properties** file. The **$mysqlRowExpirationTimeForDelete** field allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response from the server. To view/edit the current values, modify the following lines in the **.properties** file:
	`##########################`
	`# CACHE / HISTORIAN SETTINGS`
	`#`
	`collector.sl.rtview.sub=$mssqlRowExpirationTime:120`
	`collector.sl.rtview.sub=$mssqlRowExpirationTimeForDelete:0`
	
	In the example above, the **Expired** check box would be checked after 120 seconds, and the row would never be deleted. If **$mysqlRowExpirationTimeForDelete** was set to 3600, then the row would be removed from the table after 3600 seconds.

Alert Level	The current alert severity. Red indicates that one or more metrics exceeded their ALARM LEVEL threshold. Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold. Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	The total number of alerts for the server.
Connected	When checked, the server is connected.
Last Query	The status of the last query made:
Avg Exec Time	The average amount of execution time, in seconds.
Avg Process Time	The average amount of process time, in seconds.
Bytes Received	The total number of bytes received since the server was last started.
Connections	The total number of connections since the server was last started.
Delayed Writes	The total number of delayed writes.
Queries	The total number of queries.
Query Objects	The total number of query objects.
Slow Queries	The total number of slow queries.

Total Executions	The total number of executions.
Uptime	The amount of time since the server was last started, in seconds.
Concurrent	When checked, the database allows concurrent usage.
Enabled	When checked, the database is enabled for usage.
Timestamp	The data and time of the last data update.

Server Summary

View connection, performance and processing details for a single MySQL database server, as well as trending data for the number of kilobytes received and queries. Choose an instance from the **Instance** drop-down menu. Mouse over the trend graph to see performance metrics with time stamps.

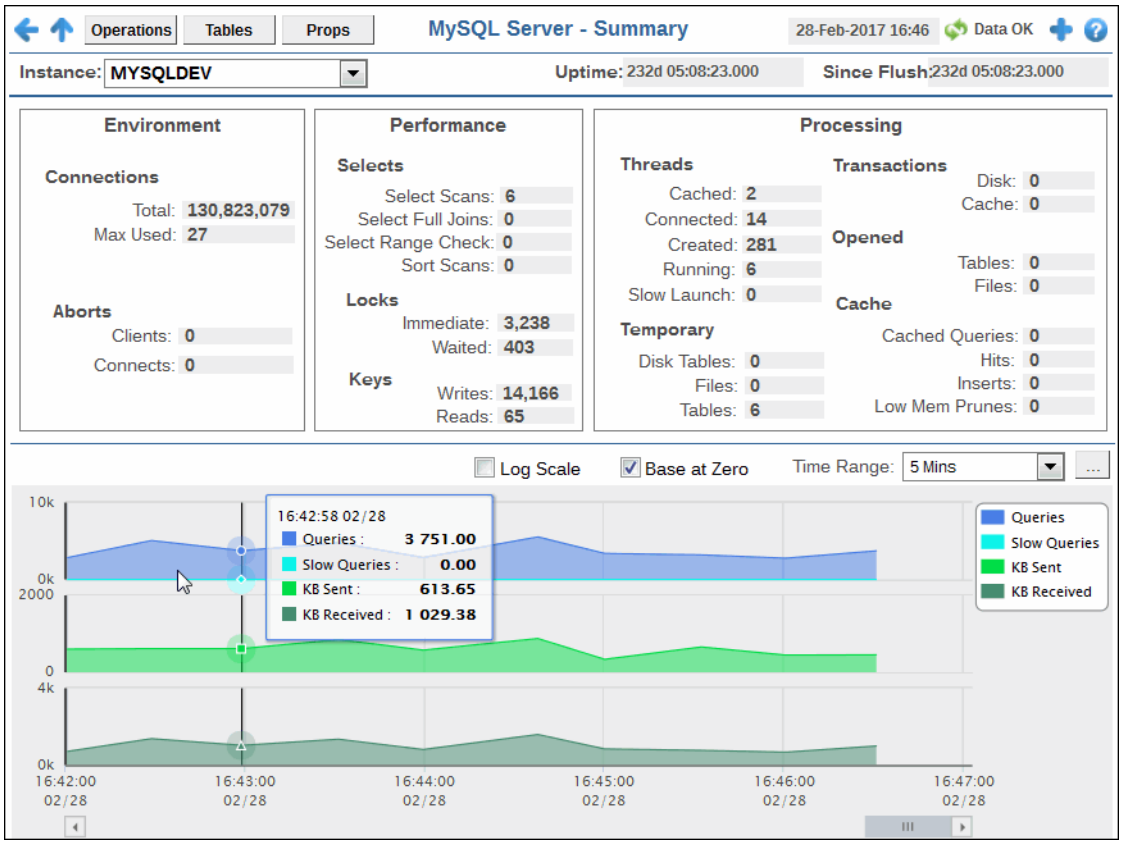

Title Bar (possible features are):

Open the previous and upper display.

Open an instance of this display in a new window.

Open the online help page for this display.

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:

Instance: Select the instance for which you want to show data in the display.

Fields and Data: For details about the data in this display, please refer to vendor documentation.

Uptime The amount of time since the server was last started, in number of days, hours, minutes and seconds.

Since Flush The amount of time since the last flush, in number of days, hours, minutes and seconds.

Performance Trend Graph Traces the following:

Queries: Traces the amount queries per second.

Slow Queries: Traces the amount of slow queries per second.

KB Sent: Traces the number of kilobytes sent per second.

KB Received: Traces the number of kilobytes received per second.

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

Base at Zero Select to use zero (**0**) as the Y axis minimum for all graph traces.

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

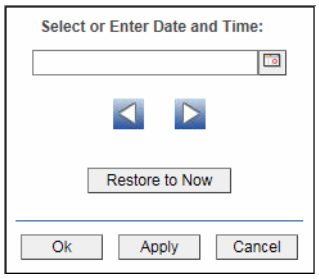

By default, the time range end point is the current time. To change the time range end point, click Calendar and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

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

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

Servers Properties

View properties and property values for a single MySQL database server.

Choose an instance from the **Instance** drop-down menu. Each table row is a different property for the selected instance. Enter a search string in the **Property Filter** field to limit the number of table rows. Click a column header to sort column data in numerical or alphabetical order.

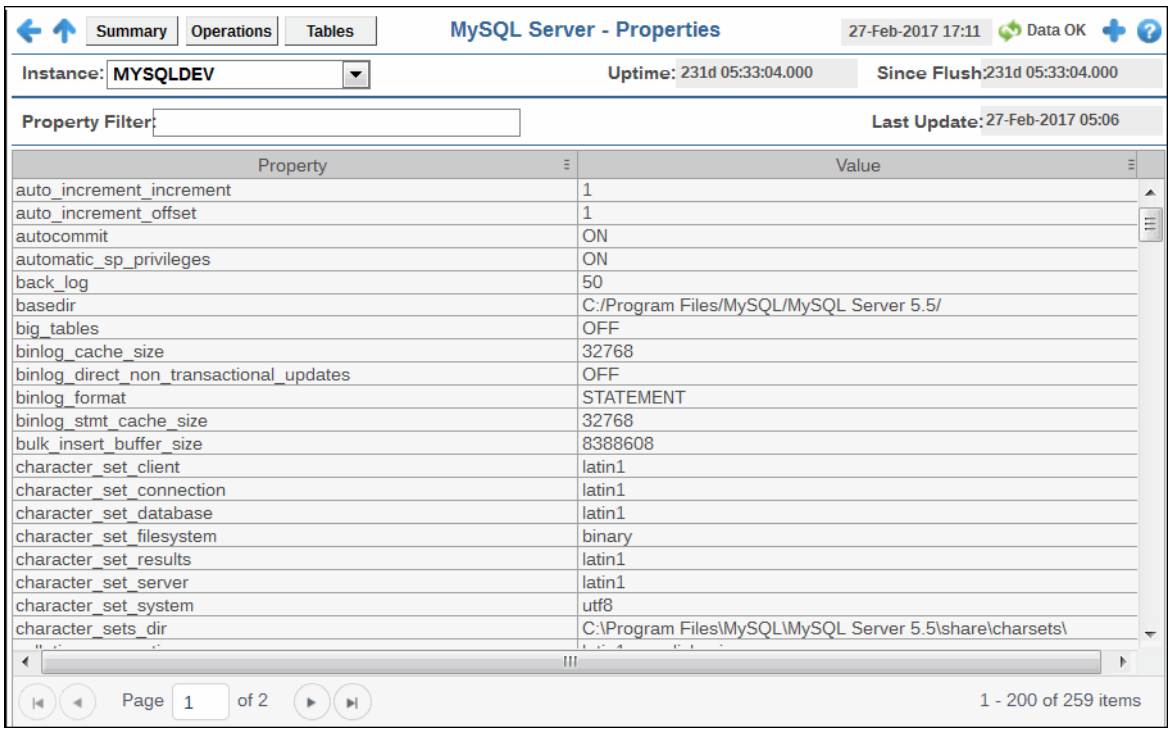

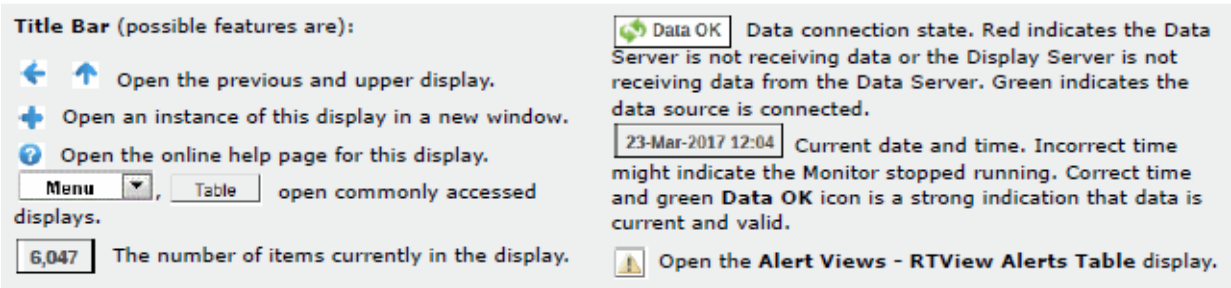

Filter By:

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

Fields and Data:

Uptime	The amount of time since the server was last started, in number of days, hours, minutes and seconds.
Property Filter:	Enter a search string to filter the number of table rows.
Since Flush	The amount of time since the last flush, in number of days, hours, minutes and seconds.

Servers Operations

View trending performance data for a single MySQL database server: **Inserts**, **Selects**, **Updates** and **Deletes**. Choose an instance from the **Instance** drop-down menu. Mouse over the trend graph to see performance metrics with time stamps.

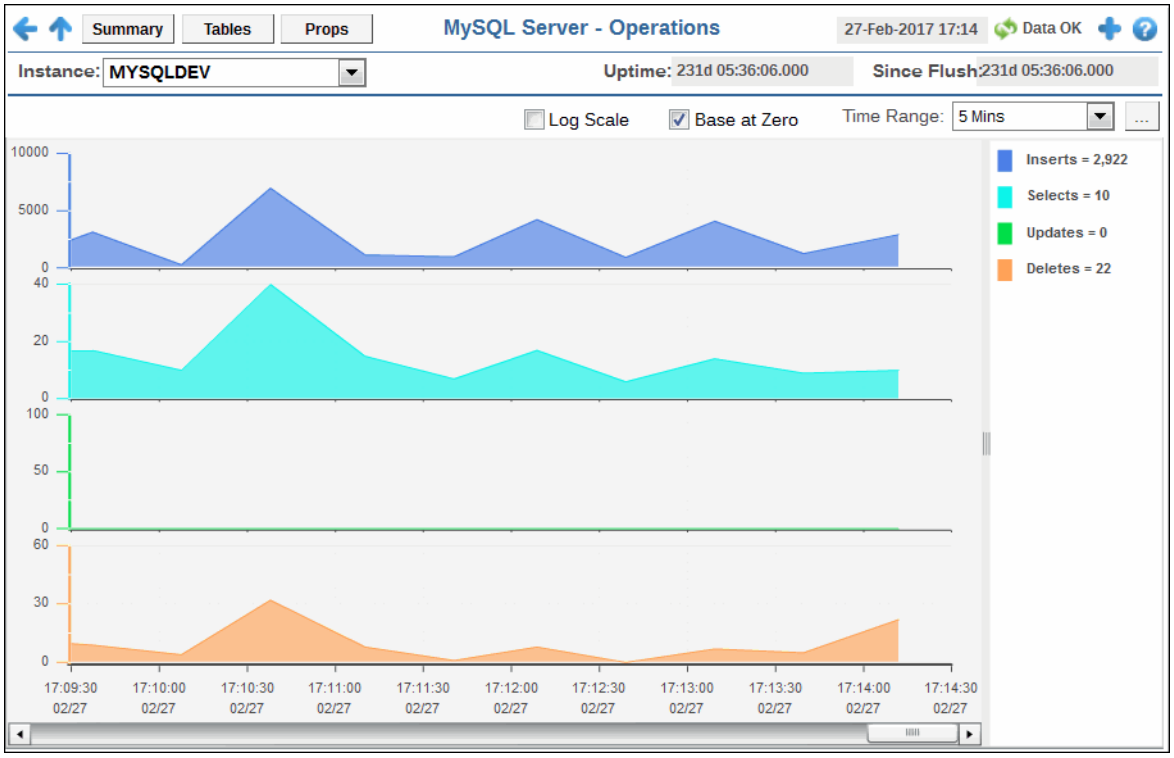

Title Bar (possible features are):

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

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

Fields and Data:

Uptime	The amount of time since the server was last started, in number of days, hours, minutes and seconds.
Property Filter:	Enter a search string to filter the number of table rows.
Since Flush	The amount of time since the last flush, in number of days, hours, minutes and seconds.

Performance Trend Graph

Traces the following:

Inserts: Traces the number of inserts per second.

Selects: Traces the number of selects per second.

Updates: Traces the number of updates per second.

Deletes: Traces the number of deletes per second.

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

Base at Zero Select to use zero (**0**) as the Y axis minimum for all graph traces.

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

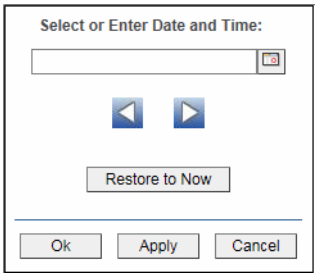

By default, the time range end point is the current time. To change the time range end point, click Calendar and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

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

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

User Tables

View performance and utilization details for cache tables for a single MySQL database server. Each row is a different cache table. Choose an instance from the **Instance** drop-down menu. Click a column header to sort column data in numerical or alphabetical order.

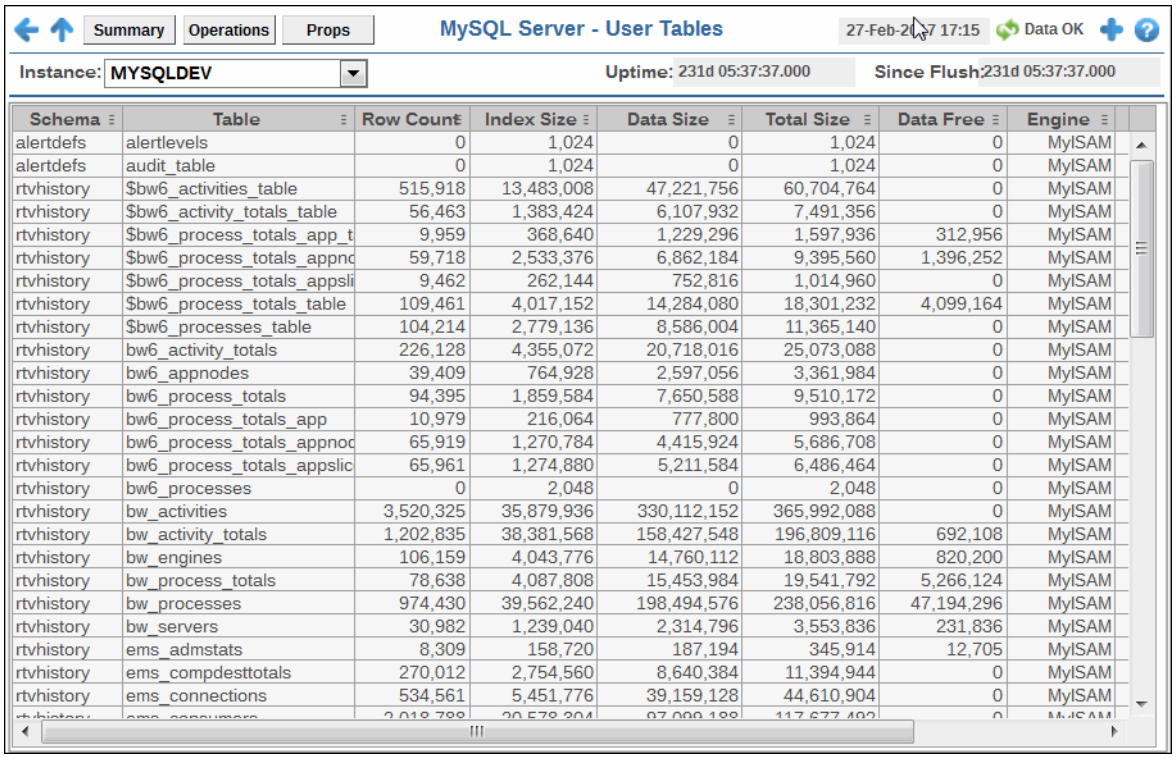

Schema	Table	Row Count	Index Size	Data Size	Total Size	Data Free	Engine
alertdefs	alertlevels	0	1,024	0	1,024	0	MyISAM
alertdefs	audit_table	0	1,024	0	1,024	0	MyISAM
rtvhistory	$bw6_activities_table	515,918	13,483,008	47,221,756	60,704,764	0	MyISAM
rtvhistory	$bw6_activity_totals_table	56,463	1,383,424	6,107,932	7,491,356	0	MyISAM
rtvhistory	$bw6_process_totals_app_t	9,959	368,640	1,229,296	1,597,936	312,956	MyISAM
rtvhistory	$bw6_process_totals_appnc	59,718	2,533,376	6,862,184	9,395,560	1,396,252	MyISAM
rtvhistory	$bw6_process_totals_appsli	9,462	262,144	752,816	1,014,960	0	MyISAM
rtvhistory	$bw6_process_totals_table	109,461	4,017,152	14,284,080	18,301,232	4,099,164	MyISAM
rtvhistory	$bw6_processes_table	104,214	2,779,136	8,586,004	11,365,140	0	MyISAM
rtvhistory	bw6_activity_totals	226,128	4,355,072	20,718,016	25,073,088	0	MyISAM
rtvhistory	bw6_appnodes	39,409	764,928	2,597,056	3,361,984	0	MyISAM
rtvhistory	bw6_process_totals	94,395	1,859,584	7,650,588	9,510,172	0	MyISAM
rtvhistory	bw6_process_totals_app	10,979	216,064	777,800	993,864	0	MyISAM
rtvhistory	bw6_process_totals_appnoc	65,919	1,270,784	4,415,924	5,686,708	0	MyISAM
rtvhistory	bw6_process_totals_appslic	65,961	1,274,880	5,211,584	6,486,464	0	MyISAM
rtvhistory	bw6_processes	0	2,048	0	2,048	0	MyISAM
rtvhistory	bw_activities	3,520,325	35,879,936	330,112,152	365,992,088	0	MyISAM
rtvhistory	bw_activity_totals	1,202,835	38,381,568	158,427,548	196,809,116	692,108	MyISAM
rtvhistory	bw_engines	106,159	4,043,776	14,760,112	18,803,888	820,200	MyISAM
rtvhistory	bw_process_totals	78,638	4,087,808	15,453,984	19,541,792	5,266,124	MyISAM
rtvhistory	bw_processes	974,430	39,562,240	198,494,576	238,056,816	47,194,296	MyISAM
rtvhistory	bw_servers	30,982	1,239,040	2,314,796	3,553,836	231,836	MyISAM
rtvhistory	ems_admstats	8,309	158,720	187,194	345,914	12,705	MyISAM
rtvhistory	ems_compdesttotals	270,012	2,754,560	8,640,384	11,394,944	0	MyISAM
rtvhistory	ems_connections	534,561	5,451,776	39,159,128	44,610,904	0	MyISAM

Title Bar (possible features are):

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:

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

Fields and Data: For details about the data in this display, please refer to vendor documentation.

Uptime The amount of time since the server was last started, in number of days, hours, minutes and seconds.

Property Filter: Enter a search string to filter the number of table rows.

Since Flush The amount of time since the last flush, in number of days, hours, minutes and seconds.

Table

Schema	The name of the database.
Table	The name of the table.
Row Count	The number of rows currently in the table.
Index Size	The size of the table indexes, in bytes.
Data Size	The size of the data stored in the table, in bytes (Total Size - Index Size = Data Size).
Total Size	The total size of the table, in bytes.
Data Free RX	The amount of available space that can be reclaimed to store new data, in bytes.
Engine	The storage engine handling the SQL operations.
Last Updated	The time of the last data update.

Docker Engines

The Docker Engines displays provide extensive visibility into the health and performance of your Docker engines. These displays are populated with performance data if you are using the RTView Monitor for Solace AMI version.

Displays are:

- "Engines Heatmap": A heatmap view of all engines and their associated metrics.
- "Engines Table": A tabular view of your engines and their associated metrics.
- "Engine Summary": Provides additional details and a way to view trending data for a single engine.
- "Containers Heatmap": A color-coded heatmap view of data for all containers for a particular host.
- "Containers Table": A tabular view of data for all containers for a particular host.
- "Container Summary": This display allows you to view current and trending data for a single container for a particular host.

Engines Heatmap

This heatmap display provides an easy-to-view interface that allows you to quickly identify the current status of each of your engines for each available metric. You can view the engines in the heatmap based on the following metrics: the current alert severity, the current alert count, the percentage of CPU used, the amount of memory used, the total incoming bytes, and the total outgoing bytes. By default, this display shows the heatmap based on the **Alert Severity** metric.

You can use the **Names** check-box ☑ to include or exclude labels in the heatmap, and you can mouse over a rectangle to see additional metrics for an engine. Clicking one of the rectangles in the heatmap opens the "Engine Summary" display, which allows you to see additional details for the selected engine.

Note: When the data for the engine being monitored expires, the color of the rectangle representing that engine in the heatmap automatically changes to a color that is not included in the color gradient bar so that you can easily identify when the data is stale. Expired data could occur for a number of reasons

including, but not limited to, the connection to the engine may have been lost, or the engine could have experienced a problem and may no longer be up-and-running.

Title Bar (possible features are):

Open the previous and upper display.

Open an instance of this display in a new window.

Open the online help page for this display.

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:

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

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

Names Select this check box to display the names of the engines at the top of each rectangle in the heatmap.

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

Auto Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value.

Note: Some metrics auto-scale automatically, even when **Auto** is not selected.

Metric Choose a metric to view in the display.

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

Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.

Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.

Green indicates that no metrics have exceeded their alert thresholds.

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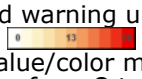 bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.

CPU Usage The percentage of CPU used by the engine. The color gradient 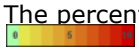 bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **DocEngineCpuUsageHigh**. The middle value in the gradient bar indicates the middle value of the range.

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

Memory The current memory usage by the engine, in kilobytes, which includes all memory regardless of when it was accessed. The color gradient bar shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of connections in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

The **Auto** option does not impact this metric.

Net Bytes In The total number of incoming bytes. The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **DocEngineNetBytesInHigh**. The middle value in the gradient bar indicates the middle value of the range.

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

Net Bytes Out The total number of outgoing bytes.The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **DocEngineNetBytesOutHigh**. The middle value in the gradient bar indicates the middle value of the range.

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

Engines Table

This table provides a view of all of your engines and their associated metric data including host, alert severity, alert count, and the current value of each gathered metric. You can click a column header to sort column data in numerical or alphabetical order, and drill-down and investigate by clicking a row to view details for the selected engine in the "Engine Summary" display

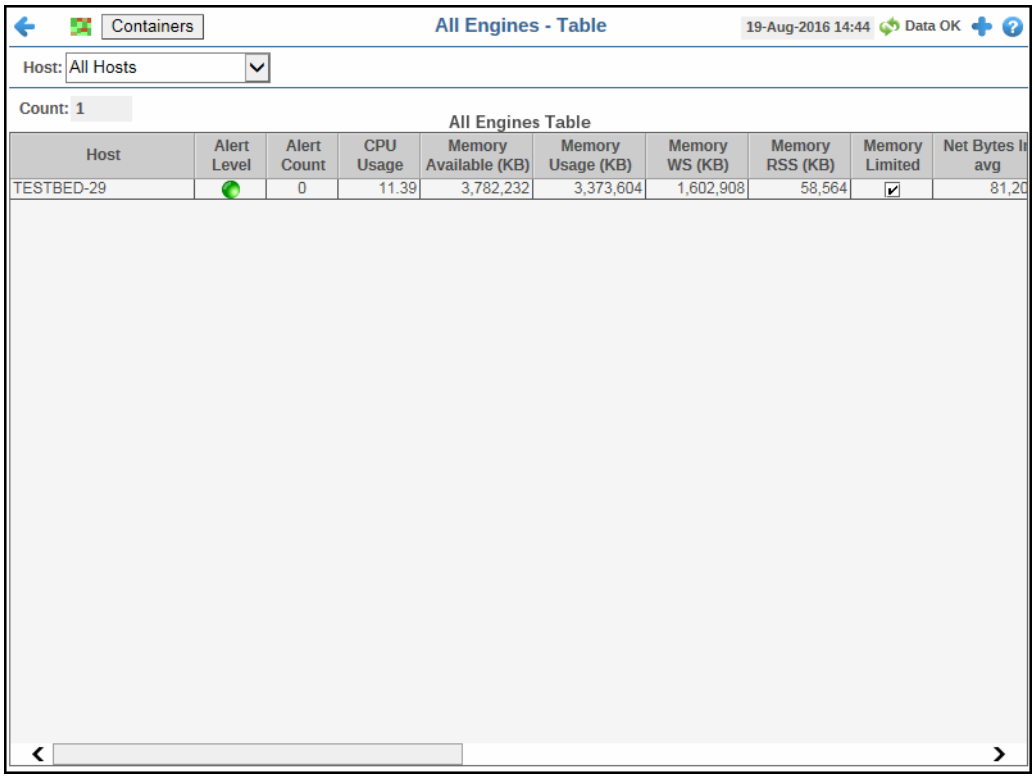

Title Bar (possible features are):

Open the previous and upper display.

Open an instance of this display in a new window.

Open the online help page for this display.

Menu, Table open commonly accessed displays.

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

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

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

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

Note: The **Containers** button takes you to "Containers Table".

Fields and Data:

Host Select the name of the host (or **All Hosts**) containing the engines for which you want to view data.

Count	The total number of engines being monitored based on your search criteria.

All Engines Table:

Host	The name of the host.
Alert Level	The current alert severity. Red indicates that one or more metrics exceeded their ALARM LEVEL threshold. Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold. Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	The total number of alerts for the host.
CPU Usage	The percentage of CPU used by the engine.
Memory Available (KB)	The amount of memory, in kilobytes, that is available to the engine.
Memory Usage (KB)	The current memory usage by the engine, in kilobytes, which includes all memory regardless of when it was accessed.
Memory WS (KB)	The amount of memory (in kilobytes) in the working set, which includes recently accessed memory, dirty memory, and kernel memory.
Memory RSS (KB)	The amount of anonymous and swap cache memory (including transparent/ hugepages), in kilobytes.
Memory Limited	When checked, the amount of memory available to the engine is limited.
Net Bytes In avg	The average number of incoming bytes per second.
Net Bytes Out avg	The average number of outgoing bytes per second.
Net Packets In avg	The average number of incoming packets per second.
Net Packets Out avg	The average number of outgoing packets per second.
Docker Version	The Docker software version of the Docker Engine.
Container OS Version	The version of the container's operating system on which the docker engine is running.
Container Kernal Version	The version of the container's Kernal in which the docker engine is running.
Expired	When checked, performance data has not been received within the time specified (in seconds) in the **Expire Time** field in the **Duration** region in the RTView Configuration Application > (**Project Name**) > **Solution Package Configuration** > **Docker** > **DATA STORAGE** tab. The **Delete Time** field (also in the **Duration** region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
Timestamp	The date and time the row data was last updated.

Engine Summary

This display allows you to view current as well as trending data for the percentage of CPU used by the engine, memory usage details, and network data details.

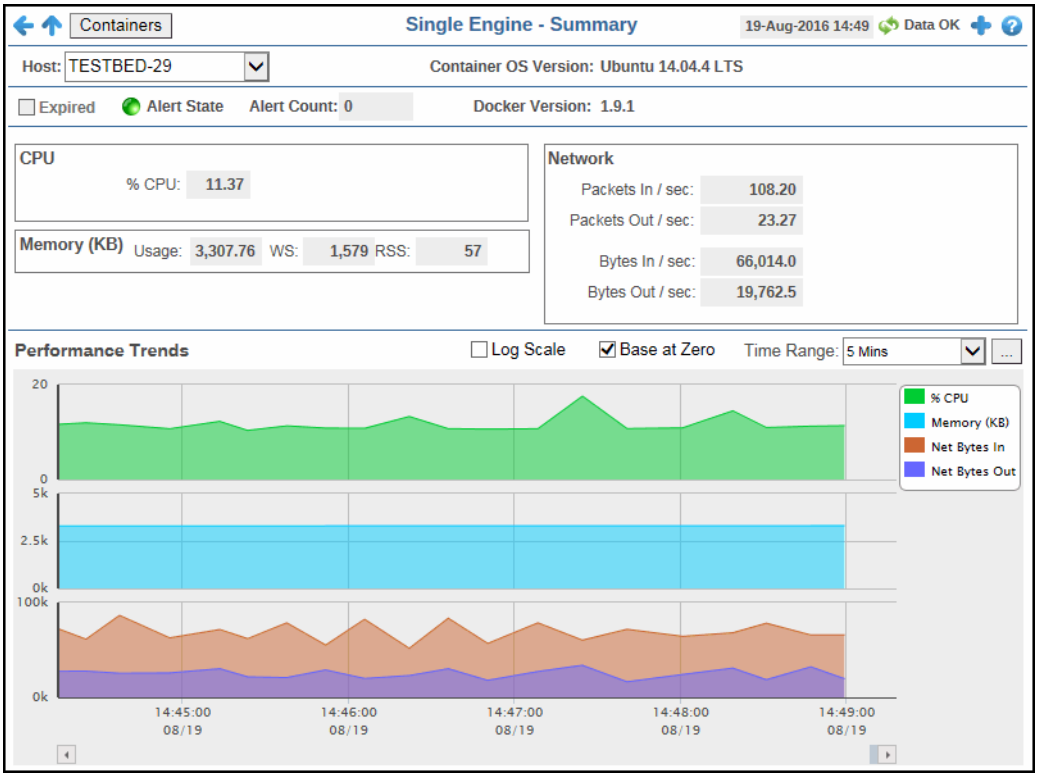

Title Bar (possible features are):

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

Note: The **Containers** button takes you to "Containers Table".

Filter By:

Host	Select the host for which you want to show data in the display.
Container OS Version	The version of the container's operating system on which the docker engine is running.

Fields and Data:

Expired When checked, performance data has not been received within the time specified (in seconds) in the **Expire Time** field in the **Duration** region in the RTView Configuration Application > (**Project Name**) > **Solution Package Configuration** > **Docker** > **DATA STORAGE** tab. The **Delete Time** field (also in the **Duration** region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.

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

Alert Count The total number of current alerts.

Docker Version The Docker software version of the Docker Engine.

CPU

% CPU The percentage of CPU used by the engine.

Memory (KB)

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

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

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

Network

Packets In/ sec The average number of incoming packets per second..

Packets Out/ sec The average number of outgoing packets per second.

Bytes In/sec The average number of incoming bytes per second.

Bytes Out/ sec The average number of outgoing bytes per second.

Performance Trends Graph Traces the following:

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

Memory (KB) -- traces the amount of memory, in kilobytes, used by the engine.

Net Bytes In -- traces the average number of incoming bytes per second.

Net Bytes Out -- traces the average number of outgoing bytes per second.

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

Base at Zero Select to use zero (**0**) as the Y axis minimum for all graph traces.

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

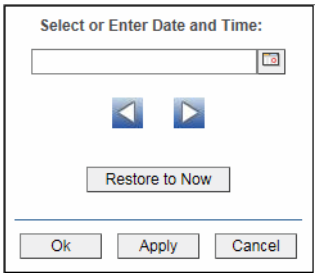

By default, the time range end point is the current time. To change the time range end point, click Calendar and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

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

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

Containers Heatmap

This heatmap display provides an easy-to-view interface that allows you to quickly identify the current status of each of your containers for each available metric. You can view the containers in the heatmap based on the following metrics: the current alert severity, the current alert count, the percentage of CPU used, and the percentage of memory used. By default, this display shows the heatmap based on the **Alert Severity** metric.

You can use the **Names** check-box to include or exclude labels in the heatmap, and you can mouse over a rectangle to see additional metrics for a container. Clicking one of the rectangles in the heatmap opens the "Container Summary" display, which allows you to see additional details for the selected container.

Note: When the data for the container being monitored expires, the color of the rectangle representing that container in the heatmap automatically changes to a color that is not included in the color gradient bar so that you can easily identify when the data is stale. Expired data could occur for a number of reasons including, but not limited to, the connection to the container may have been lost, or the container could have experienced a problem and may no longer be up-and-running.

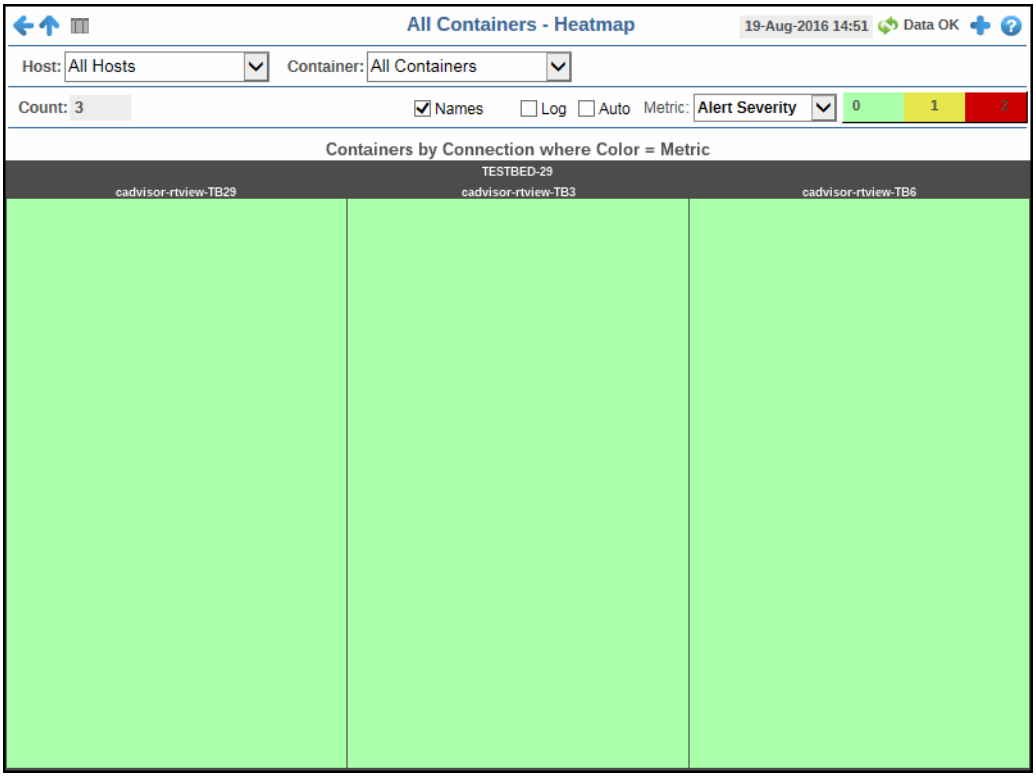

Title Bar (possible features are):

Open the previous and upper display.

Open an instance of this display in a new window.

Open the online help page for this display.

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:

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

Container Select the container (or **All Containers**) for which you want to show data in the heatmap..

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

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

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

Auto Select to enable auto-scaling. When auto-scaling is activated, the color gradient bar's maximum range displays the highest value.

Note: Some metrics auto-scale automatically, even when **Auto** is not selected.

Metric Choose a metric to view in the display.

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

Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.

Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.

Green indicates that no metrics have exceeded their alert thresholds.

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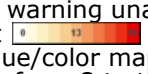 bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.

CPU Usage The percentage of CPU used by the container. The color gradient 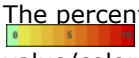 bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **DocContainerCpuUsageHigh**. The middle value in the gradient bar indicates the middle value of the range.

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

Memory The current memory usage by the container, in kilobytes, which includes all memory regardless of when it was accessed. The color gradient bar 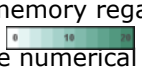 shows the range of the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of connections in the heatmap. The middle value in the gradient bar indicates the middle value of the range.

The **Auto** option does not impact this metric.

Net Bytes In The number of incoming bytes per second. The color gradient 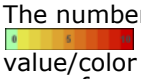 bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **DocContainerNetBytesInHigh**. The middle value in the gradient bar indicates the middle value of the range.

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

Net Bytes Out The number of outgoing bytes per second.The color gradient 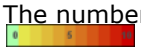 bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **DocContainerNetBytesOutHigh**. The middle value in the gradient bar indicates the middle value of the range.

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

Containers Table

This display allows you to view details in a table format for one container on a particular host, for all containers on a particular host, for a particular container on all hosts, or for all containers on all hosts. You can drill-down and view the details for a particular container in the "Container Summary" display by clicking on a row in the resulting table.

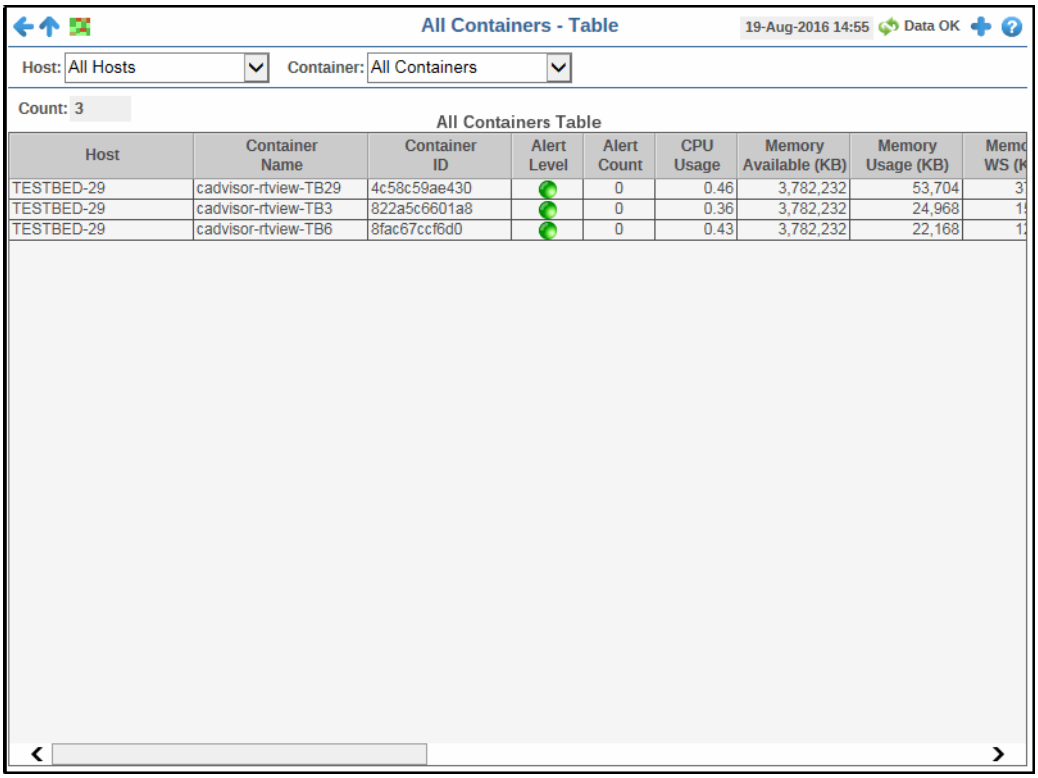

Title Bar (possible features are):

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 includes these filtering options:

Host	Select the host for which you want to show data in the display.
Container	Select the container (or **All Containers**) for which you want to view data..
Count	Lists the total number of containers (rows) found using the search parameters.

All Containers Table

Host	The name of the host.
Container Name	The name of the container.
Container ID	The absolute container name.
Alert Level	The current alert status. Red indicates that one or more metrics exceeded their ALARM LEVEL threshold. Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold. Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	Total number of alerts for the process.
CPU Usage	The percentage of CPU used by the container.
Memory Available (KB)	The amount of memory, in kilobytes, that is available to the container.
Memory Usage (KB)	Current memory usage by the container, in kilobytes, which includes all memory regardless of when it was accessed.
Memory WS (KB)	The amount of memory (in kilobytes) in the working set, which includes recently accessed memory, dirty memory, and kernel memory.
Memory RSS (KB)	The Resident Set Size, which is the amount of anonymous and swap cache memory (including transparent/hugepages), in kilobytes.
Memory Limited	When checked, the amount of memory available to the container is limited. If not checked, then the amount of memory available to the container is unlimited, which means the amount of memory available to the container is the same as the memory available to the engine.
Net Bytes In avg	The average number of incoming bytes per second.
Net Bytes Out avg	The average number of outgoing bytes per second.
Net Packets In avg	The average number of incoming packets per second.
Net Packets Out avg	The average number of outgoing packets per second.
Uptime	The amount of time (in seconds) that the container has been up and running.
Running	When checked, this check box indicates that the container is running.
Status	The current status of the container. Values are: **Up** - indicates that the container is up and running, and lists the amount of time the container has been up and running (**Uptime**). **Created** - indicates that the container has been created but is currently not in use. **Exited** - indicates that the container has been stopped, and lists the error code as well as the amount of time since the container was stopped.
Starts	The number of times the container (re)started within the time specified (in seconds) in the **$docEventCacheTimeRange** field in the **conf\rtvapm_dockermon.properties** file. The default is 3600 seconds (1 hour). For example, by default, this row column lists the number of times the container has (re)started in the past hour. This number provides a good indication of the stability of the container; the higher the number, the more unstable the container.

Expired	When checked, performance data has not been received within the time specified (in seconds) in the **Expire Time** field in the **Duration** region in the RTView Configuration Application > (**Project Name**) > **Solution Package Configuration** > **Docker** > **DATA STORAGE** tab. The **Delete Time** field (also in the **Duration** region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.
Timestamp	The date and time the row data was last updated.

Container Summary

This display provides a view of the current and historical metrics for a single container. You can view the current information pertaining to CPU usage percentage, Memory details, Disk read and write details, and network data details in the upper portion of the display. The trend graph in the bottom half of the display traces the current and historical CPU usage, the average memory used, and the number of incoming and outgoing network bytes.

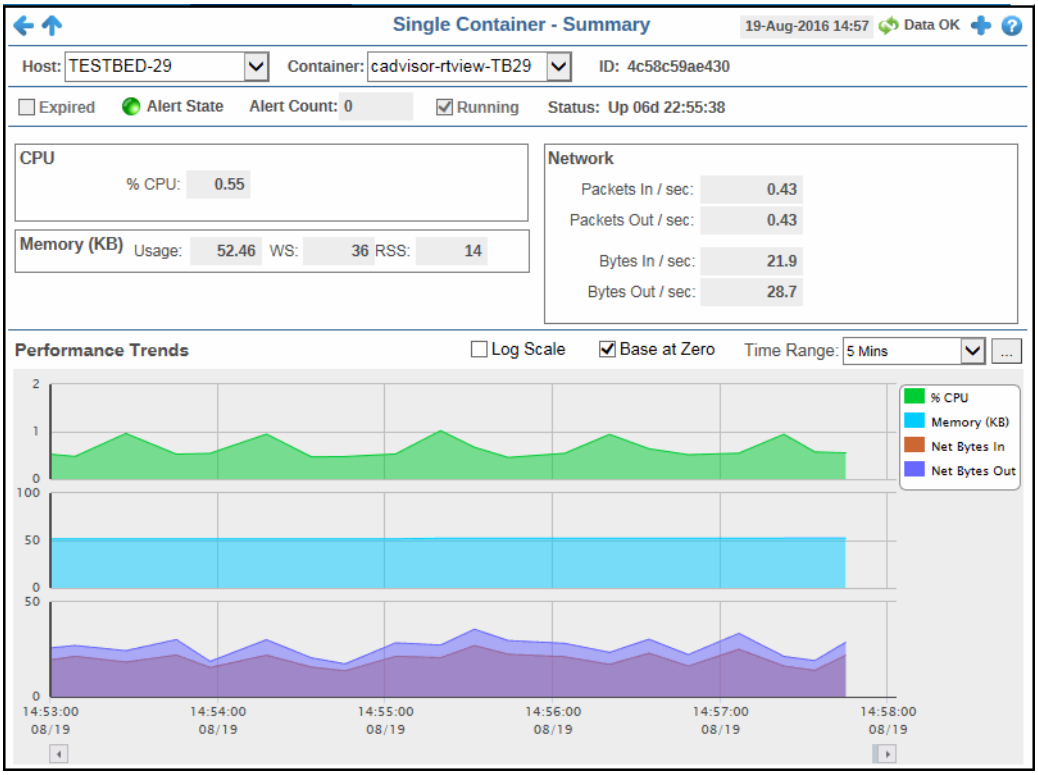

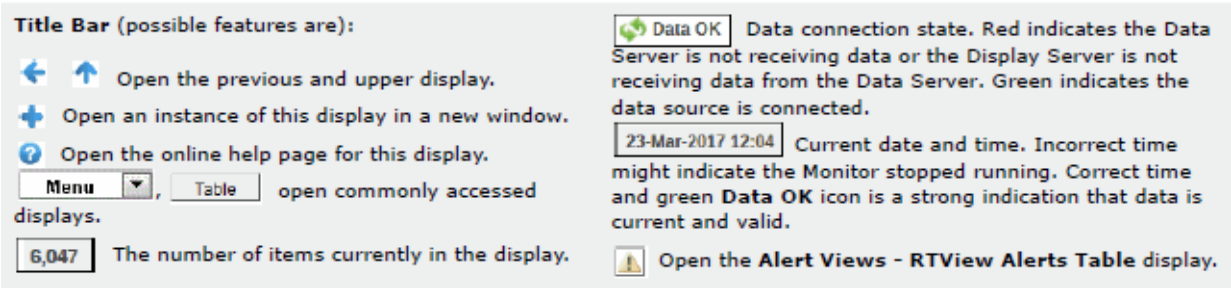

Filter By:

The display might include these filtering options:

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

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

ID The absolute container name.

Fields and Data:

Expired When checked, performance data has not been received within the time specified (in seconds) in the **Expire Time** field in the **Duration** region in the RTView Configuration Application > (**Project Name**) > **Solution Package Configuration** > **Docker** > **DATA STORAGE** tab. The **Delete Time** field (also in the **Duration** region) allows you to define the amount of time (in seconds) in which the row will be removed from the table if there is no response.

Alert State The current alert severity.

Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.

Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.

Green indicates that no metrics have exceeded their alert thresholds.

Alert Count The total number of current alerts.

Running When checked, this check box indicates that the container is running.

Status The current status of the container. Values are:

Up - indicates that the container is up and running, and lists the amount of time the container has been up and running (**Uptime**).

Created - indicates that the container has been created but is currently not in use.

Exited - indicates that the container has been stopped, and lists the error code as well as the amount of time since the container was stopped.

CPU

% CPU The percentage of CPU used by the container.

Memory (KB)

Usage The current memory usage by the container, in kilobytes, which includes all memory regardless of when it was accessed.

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

	RSS	The Resident Set Size, which is the amount of anonymous and swap cache memory (including transparent/hugepages), in kilobytes.
Network		
	Packets In/ sec	The average number of incoming packets per second.
	Packets Out/ sec	The average number of outgoing packets per second.
	Bytes In/sec	The average number of incoming bytes per second.
	Bytes Out/sec	The average number of outgoing bytes per second.

Performance Trends Graph

Traces the following:

% CPU -- traces percentage of CPU used by the container.

Memory (KB) -- traces the current memory usage by the container, in kilobytes, which includes all memory regardless of when it was accessed.

Net Bytes In -- traces the average number of incoming bytes per second.

Net Bytes Out -- traces the average number of outgoing bytes per second.

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

Base at Zero — Select to use zero (**0**) as the Y axis minimum for all graph traces.

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

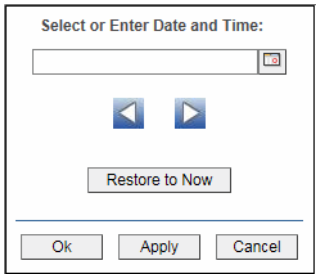

By default, the time range end point is the current time. To change the time range end point, click Calendar and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

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

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

Hosts

Hosts displays provide extensive visibility into the health and performance of your hosts.

Displays are:

- "All Hosts Heatmap"
- "All Hosts Table"
- "All Hosts Grid"
- "All Processes Table"
- "All Network Table"
- "All Storage Table"
- "Host Summary"

All Hosts Heatmap

View the most critical alert states pertaining to your hosts. Use this display to quickly identify hosts with critical alerts.

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

Choose a domain or **All Domains** from the **Domain** drop-down menu to filter data shown in the display. Choose a different metric to display from the **Metric** drop-down menu. Mouse over a rectangle to see additional metrics. By default, this display shows **Alert Severity**.

Drill-down and investigate a host by clicking a rectangle in the heatmap to view details in the **Host Summary** display.

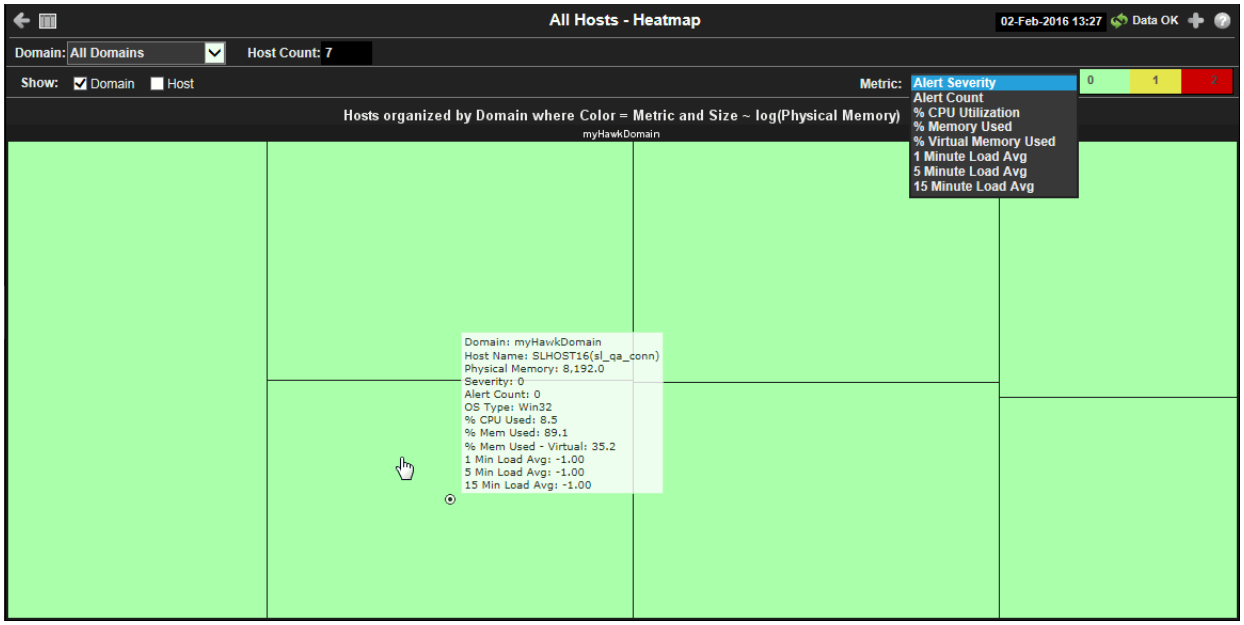

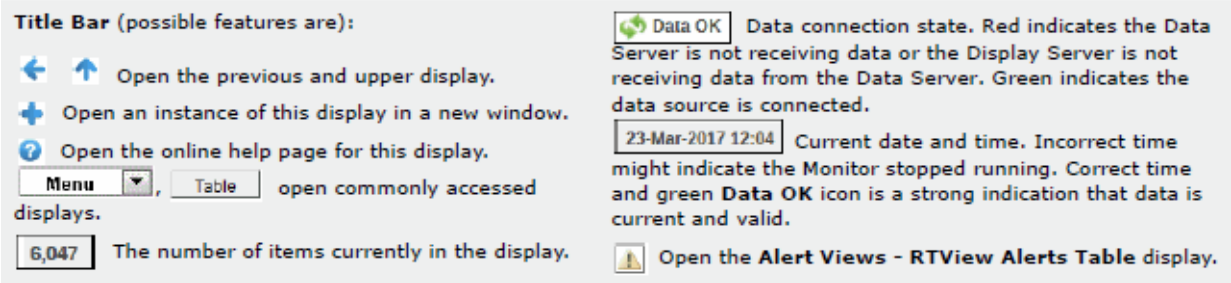

Filter By:
The display might include these filtering options:

Domain: Choose a domain to show data for in the display. Domain names are specified when your administrator configures your Data Server to collect Hawk data, and applies to all host data collected from Hawk by that Data Server.

Fields and Data:

Host Count: The total number of hosts currently shown in the display.

Show:

Domain When selected, includes the Domain name in the display.

Host When selected, includes the Host name in the display.

Metric Choose a metric to view in the display.

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

Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.

Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.

Green indicates that no metrics have exceeded their alert thresholds.

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

% CPU Utilization The percent of CPU used in the heatmap rectangle. The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

% Memory Used The percent of memory used in the heatmap rectangle. The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

% Virtual Memory Used The percent of virtual memory used in the heatmap rectangle. The color gradient bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from **0** to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average count.

1 Minute Load Avg	The average number of processes running over 1 minute.
5 Minute Load Avg	The average number of processes running over 5 minutes.
15 Minute Load Avg	The average number of processes running over 15 minutes.

All Hosts Table

View host utilization data in a tabular format. Use this display to see all available data for this View.

Each row in the table is a different host. Choose a domain or **All Domains** from the **Domain** drop-down menu. Click a column header to sort column data in numerical or alphabetical order. Drill-down and investigate by clicking a row to view details for the selected application in the **Host Summary** display.

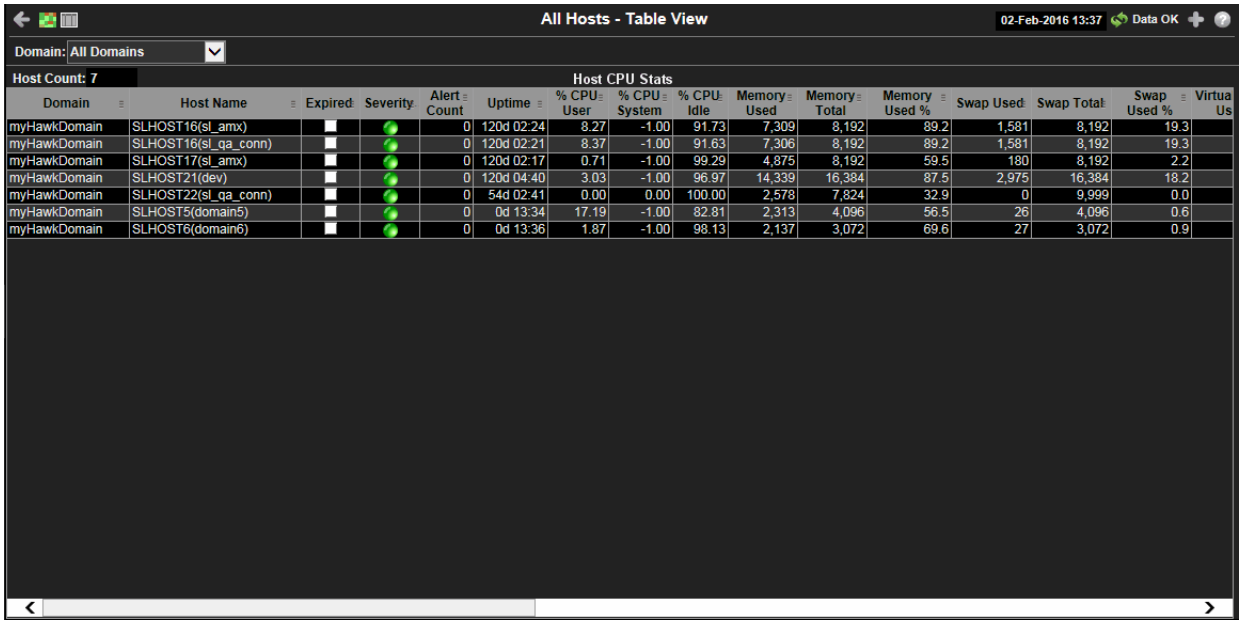

Title Bar (possible features are):

Open the previous and upper display.

Open an instance of this display in a new window.

Open the online help page for this display.

Menu, Table open commonly accessed displays.

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

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

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

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

Filter By:
The display might include these filtering options:

Domain:	Choose a domain to show data for in the display.

Fields and Data:

Host Count: The total number of hosts in the table.

Table:
Each row in the table is a different host.

Domain	The domain in which the host resides. Domain names are specified when your administrator configures your Data Server to collect Hawk data, and applies to all host data collected from Hawk by that Data Server.
Host Name	The name of the host.
Expired	When checked, data has not been received from this host in the specified amount of time. The host will be removed from the Monitor in the specified amount of time. The default setting is **60** seconds.
Severity	The maximum level of alerts in the row. Values range from **0** - **2**, as indicated in the color gradient bar, where **2** is the highest Alert Severity: Red indicates that one or more metrics exceeded their ALARM LEVEL threshold. Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold. Green indicates that no metrics exceeded their alert thresholds.
Alert Count	The total number of active alerts associated with the host.
Uptime	The amount of time the application has been running, in the following format: **0d 00:00 <days>d <hours>:<minutes>:<seconds>** For example: **10d 08:41:38**
% CPU Used	The amount of CPU used, in percent.
% CPU System	The amount of CPU used, in percent.
% CPU Idle	The amount of CPU not used, in percent.
Memory Used	The amount of memory, in megabytes, currently used.
Memory Total	The total amount of memory, in megabytes.
Memory Used%	The amount of memory used, in percent.
Swap Used	The amount of swap space, in megabytes, currently used.
Swap Total	The total amount of swap space, in megabytes.
Swap Used %	The amount of swap space used, in percent.
Virtual Mem(ory) Used	The amount of virtual memory currently used, in megabytes.
Virtual Mem(ory) Total	The total amount of virtual memory, in megabytes.
Virtual Mem(ory) Used%	The amount of virtual memory used, in percent.
Load Avg 1 Minute	The average number of processes running over 1 minute.

Load Avg 5 Minute	The average number of processes running over 5 minutes.
Load Avg 15 Minute	The average number of processes running over 15 minutes.
OS Type	The type of operating system (for example, Linux, HP-UX, Windows 2003).
OS Description	The name of the operating system.
OS Version	The operating system version.
CPU Model	The CPU model.
# CPUs	The number of node connections.
Agent Type	The type of agent from which the data was collected: **HOSTMON** (a SL Host Agent), **Hawk**, **WMI** or **SNMP**.
Agent Class	The specific version of the agent software.
Source	The name of the SL Data Server where the host data was collected.
Timestamp	The date and time the data was last updated.

All Hosts Grid

This grid provides a list view of utilization metrics for all hosts. Use this display to track and view in parallel the general performance of your hosts. Drill down and investigate by clicking a host to view details in the **Host Summary** display.

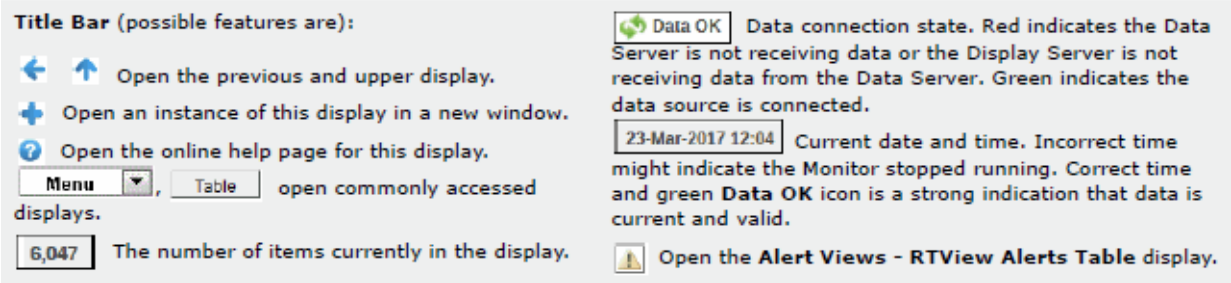

Filter By:
The display might include these filtering options:

Domain:	Choose a domain to show data for in the display. Domain names are specified when your administrator configures your Data Server to collect Hawk data, and applies to all host data collected from Hawk by that Data Server.
Host Count	Displays the number of hosts (including expired hosts) listed in the display.
Time Range:	Choose a time range to show data for in the display. Options are: **All Data**, **2 Mins**, **5 Mins**, **20 Mins**, **1 Hour, 2 Hours**, **4 Hours**, **8 Hours**, **24 Hours**, **2 Days and 7 Days**.

Grid
Utilization data shown for hosts in the selected domain.

Host Name	The name of the host.	
OS Type	The name of the operating system.	
Uptime	The amount of time (days, hours, seconds) the operating system has been running.	
Phys Mem	The amount of physical memory used, in megabytes.	
Virtual Mem	The amount of virtual memory used, in megabytes.	
Load Avg	**1**	The average number of processes running over 1 minute.
	5	The average number of processes running over 5 minutes.
	15	The average number of processes running over 15 minutes.
CPU Usage	The bar graph shows the amount of CPU currently used.	
VMem Usage	The bar graph shows the amount of virtual memory currently used.	
Trend Graphs		
	CPU	Traces the amount of CPU currently used.
	VM Usage	Traces the amount of virtual memory currently used.
	Rx KB/s	Traces the amount data currently being received per second.
	Tx KB/s	Traces the amount data currently being transmitted per second.

All Processes Table

View host utilization data in a tabular format. Use this display to see all available data for this View. Each row in the table is a different host. Choose a domain or **All Domains** and a host or **All Hosts** from the drop-down menus. Click a column header to sort column data in numerical or alphabetical order. Drill-down and investigate by clicking a row to view details for the selected application in the **Host Summary** display.

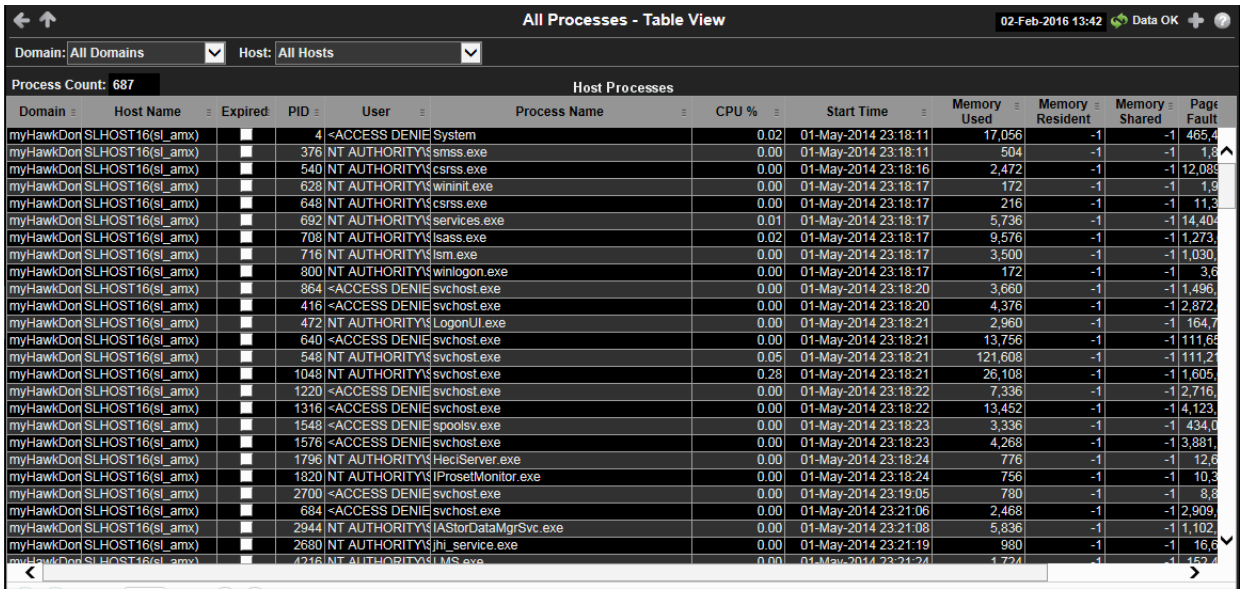

Title Bar (possible features are):

Open the previous and upper display.

Open an instance of this display in a new window.

Open the online help page for this display.

Menu, Table open commonly accessed displays.

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

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

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

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

Filter By:
The display might include these filtering options:

Domain: Choose a domain to show data for in the display. Domain names are specified when your administrator configures your Data Server to collect Hawk data, and applies to all host data collected from Hawk by that Data Server.

Host: Choose a host to show data for in the display.

Fields and Data:

Process Count: The total number of processes in the table.

Table:
Each row in the table is a different host.

Domain The domain in which the host resides.

Host Name	The name of the host.
Expired	When checked, data has not been received from this host in the specified amount of time. The host will be removed from the Monitor in the specified amount of time. The default setting is 60 seconds.
PID	The process ID.
User	The user name.
Process Name	The name of the process.
CPU%	The amount of CPU used, in percent.
Start Time	The host start time, in the following format: **0d 00:00 <days>d <hours>:<minutes>:<seconds>** For example: **10d 08:41:38**
Memory Used	The amount of memory currently used, in megabytes.
Memory Resident	The amount of memory currently used by the process that resides in physical memory and is not paged out. Set to **-1** when the data is not available from an agent. (Hawk does not provide this data.)
Memory Shared	The amount of physical memory that is shared with other processes. Set to **-1** when the data is not available from an agent. (Hawk does not provide this data.)
Page Faults	The number of page faults.
Page Faults /sec	The number of page faults per second.
Timestamp	The date and time the data was last updated.

All Network Table

View network interface data in a tabular format. Each row in the table is a different network interface card (NIC). Choose a domain or **All Domains** and a host or **All Hosts** from the drop-down menus. Click a column header to sort column data in numerical or alphabetical order.

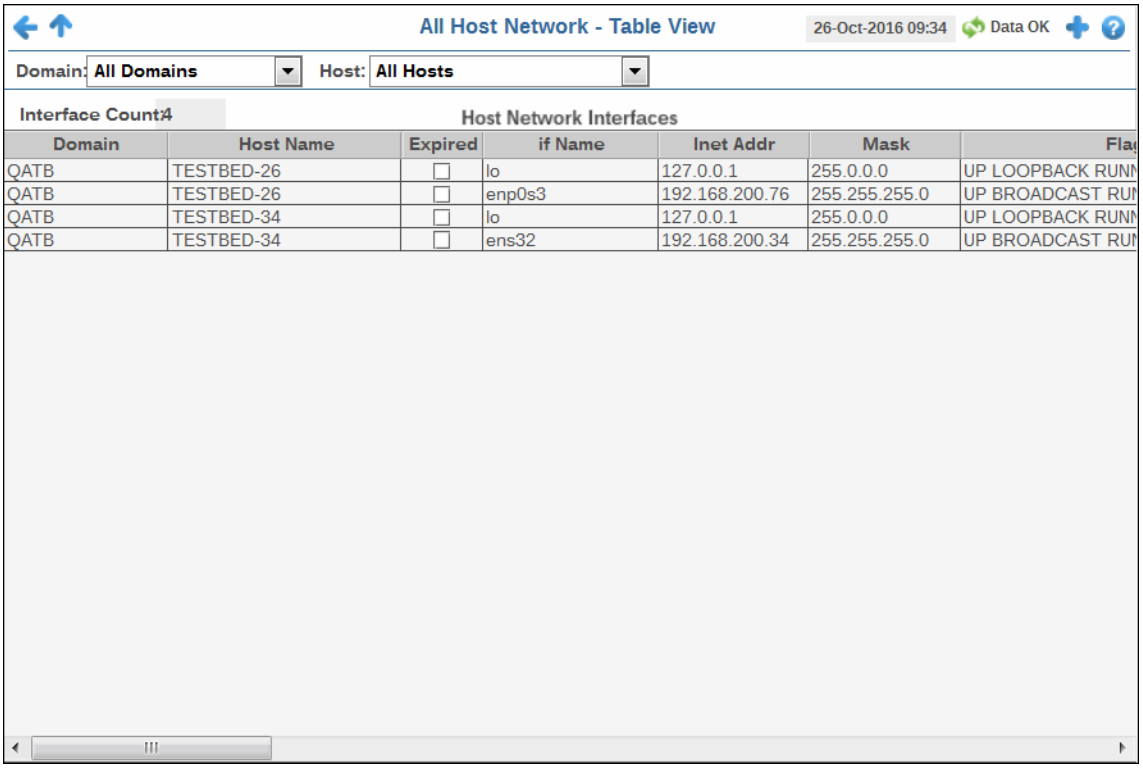

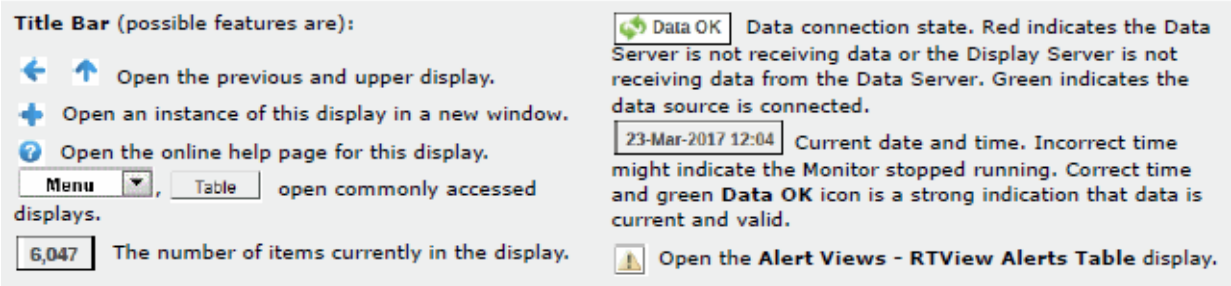

.

Filter By:
The display might include these filtering options:

Domain: Choose a domain for which to show NIC data. Domain names are specified when your administrator configures your Data Server.

Host: Choose a host for which to show NIC data.

Fields and Data:

Interface Count: The total number of NICs in the table.

Table:
Each row in the table is a different NIC.

Domain	The domain in which the NIC resides.
Host Name	The name of the NIC in which the network interface resides.
Expired	When checked, data has not been received from this NIC in the specified amount of time. The NIC will be removed from the Monitor in the specified amount of time. The default setting is 60 seconds.
if Name	The name of the NIC.
Inet Addr	The NIC IP address.
Mask	The NIC subnet mask IP address.
Flags	Descriptive text for NIC flag.
MTU	The largest size packet or frame for the NIC.
Metric	Indicates...
Point To Point	Indicates whether the NIC is a point to point configuration.
Broadcast	Indicates whether the NIC is a broadcast configuration.
rxKBytes	The total number of kilobytes received by the NIC.
rxPackets	The total number of packets received by the NIC.
rxDropped	The total number of received packets that were dropped by the NIC.
rxErrors	The total number of received errors on the NIC.
rxOverruns	The total number of received overruns on the NIC.
rxFrame	The total number of received frames on the NIC.
txKBytes	The total number of kilobytes transmitted by the NIC.
txPackets	The total number of packets transmitted by the NIC.
txDropped	The total number of transmitted packets that were dropped by the NIC.
txErrors	The total number of transmission errors for the NIC.
txOverruns	The total number of transmission overruns for the NIC.
txCollisions	The total number of transmission collisions for the NIC.
txCarrier	The total number of transmission carrier errors for the NIC.
MAC Address	The NIC MAC address.
Rx KB/s	The number of kilobytes received per second.
Tx KB/s	The number of kilobytes transmitted per second.
Rx Packets/s	The number of packets received per second.

Tx Packets/s The number of packets transmitted per second.

Timestamp The date and time the data was last updated.

All Storage Table

View storage data in a tabular format. Each row in the table is a different storage partition. Choose a domain or **All Domains** and a host or **All Hosts** from the drop-down menus. Click a column header to sort column data in numerical or alphabetical order.

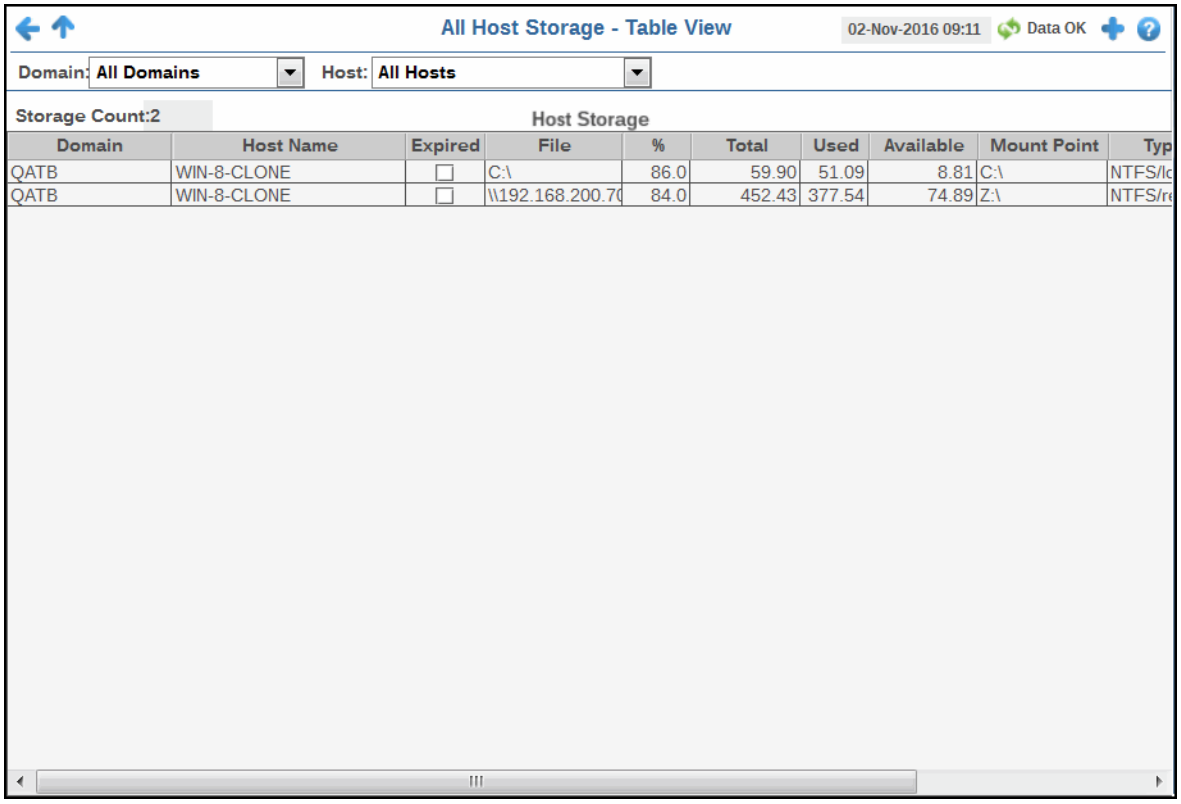

Title Bar (possible features are):

Open the previous and upper display.

Open an instance of this display in a new window.

Open the online help page for this display.

Menu, Table open commonly accessed displays.

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

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

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

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

.

Filter By:
The display might include these filtering options:

Domain:	Choose a domain or **All Domains** to show data for in the display. Domain names are specified when your administrator configures your Data Server to collect Hawk data, and applies to all host data collected from Hawk by that Data Server.
Host:	Choose a host or **All Hosts** to show data for in the display.

Fields and Data:

Storage Count:	The total number of storage partitions in the table.

Table:
Each row in the table is a different host.

Domain	The domain in which the host resides.
Host Name	The name of the host in which the storage partition resides.
Expired	When checked, data has not been received from this host in the specified amount of time. The host will be removed from the Monitor in the specified amount of time. The default setting is 60 seconds.
File System	The storage partition location.
% Used	The amount of storage partition used, in percent.
Total Size (GB)	The storage partition size, in gigabytes.
Used (GB)	The amount of storage partition used, in gigabytes.
Available (GB)	The amount of storage partition available, in gigabytes.
Mount Point	The storage partition parent directory.
Type	The file system type.
Timestamp	The date and time the data was last updated.

Host Summary

This display provides a detailed view of utilization metrics for a single server.

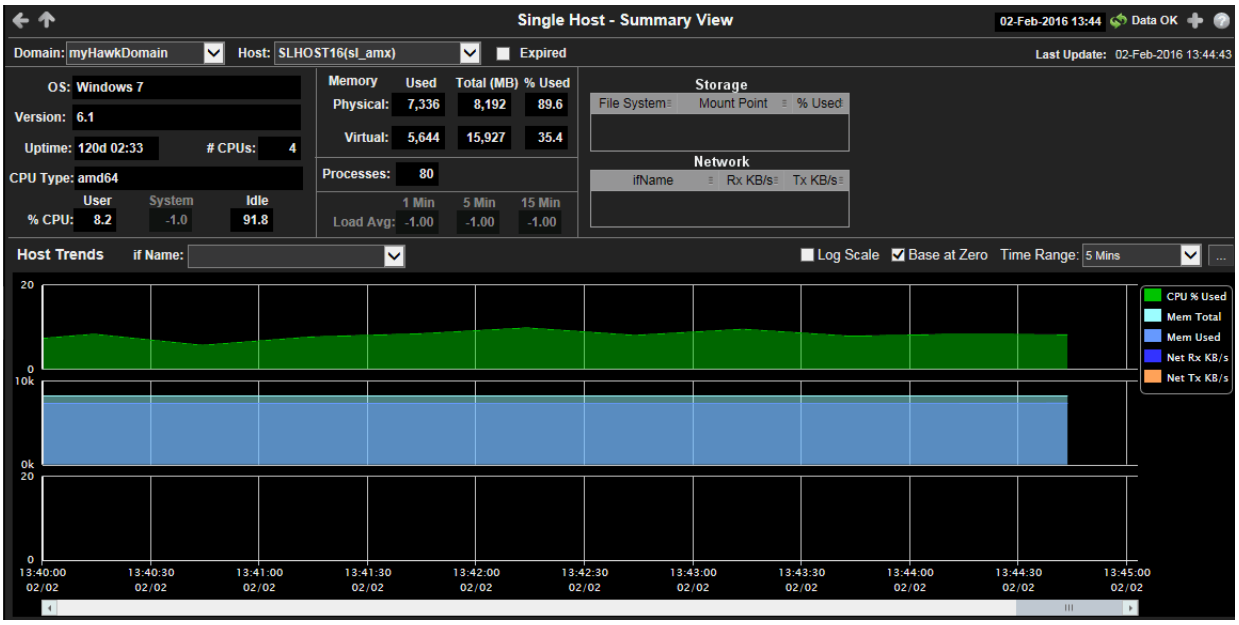

Title Bar (possible features are):

Open the previous and upper display.

Open an instance of this display in a new window.

Open the online help page for this display.

Menu, Table open commonly accessed displays.

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

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

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

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

Filter By:
The display might include these filtering options:

Domain:	Choose a domain to show data for in the display. Domain names are specified when your administrator configures your Data Server to collect Hawk data, and applies to all host data collected from Hawk by that Data Server.
Host:	Choose a host to show data for in the display.
Expired	When checked, data has not been received from this host in the specified amount of time. The host will be removed from the Monitor in the specified amount of time. The default setting is **60** seconds.
Last Update	The time the display was last updated.

Fields and Data:
Data describes the selected host except where noted.

OS:	The operating system.
Version:	The operating system version.
Uptime:	The number of days, hours and minutes since started.

	#CPUs	The number of node connections.
CPU Type:	The type of CPU.	
%CPU	**User**	The amount of CPU used by the user, in percent.
	System	The amount of CPU used by the system, in percent.
	Idle	The amount of CPU that is not used, in percent.
Physical Memory	**Used**	The amount of physical memory used, in kilobytes.
	Total(MB)	The amount of physical memory available, in kilobytes.
	%Used	The amount of physical memory used, in percent.
Virtual Memory	**Used**	The amount of virtual memory used, in kilobytes.
	Total(MB)	The amount of virtual memory available, in kilobytes.
	%Used	The amount of virtual memory used, in percent.
Processes	The number of processes running.	
Load Avg:	**1 Min**	The average number of processes running over 1 minute.
	5 Min	The average number of processes running over 5 minutes.
	15 Min	The average number of processes running over 15 minutes.
Storage	**File System**	The amount of storage space used for the file system, in kilobytes.
	Mount Point	The name used by the operating system to mount and provide an entry point to other storage volumes.
	%Used	The amount of storage space used, in percent.
Network	**ifName**	The name assigned to the network interface by the operating system.
	RxKB/s	The amount of network data received per second, in kilobytes.
	TxKB/s	The amount of network data transmitted per second, in kilobytes.

Trend Graphs
Traces metrics for the selected host.

- **CPU% Used**: The amount of CPU used, in percent.
- **Mem Total:** The amount of available memory, in kilobytes.
- **Mem Used**: The amount of memory used, in kilobytes.
- **Net Rx KB/s**: The amount of network data received per second, in kilobytes.
- **Net Tx KB/s**: The amount of network data transmitted per second, in kilobytes.

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

Base at Zero Select to use zero (**0**) as the Y axis minimum for all graph traces.

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

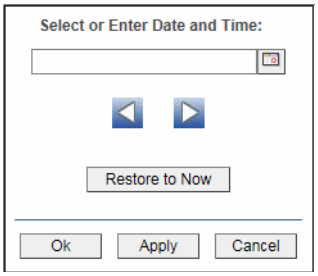

By default, the time range end point is the current time. To change the time range end point, click Calendar ... and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

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

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

Alert Views

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

- "Alert Detail Table": Shows current alert data. Use this time-ordered tabular view to track, manage and assign alerts.

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.

The color coded navigation tree shows the contents of the CMDB hierarchically ordered. Choose a node to filter alerts shown in the table. The **Alerts Table** only shows alerts associated with the node you select. A green indicator means the node has no associated alerts. A red indicator means the node has one or more associated alerts.

Service name labels are appended with the Environment and number of alerts. For example, the following illustrates that the **TBE** Service currently has no (**0**) associated alerts in the **PRODUCTION** Environment.

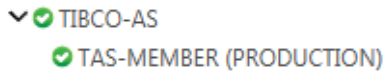

Each row in the table is a different active alert. Select one or more rows, right-click and choose **Alert** to see all actions that you can perform on the selected alert(s). Choose **Alert / Set Filter Field** to apply the selected cell data to the **Field Filter** and **Search Text** fields. Or enter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

Click a column heading to sort the table on that column data.

Optionally, you can use the **$rtvUserShowDualTables** substitution to add a table that lists alerts owned by the logged in user.

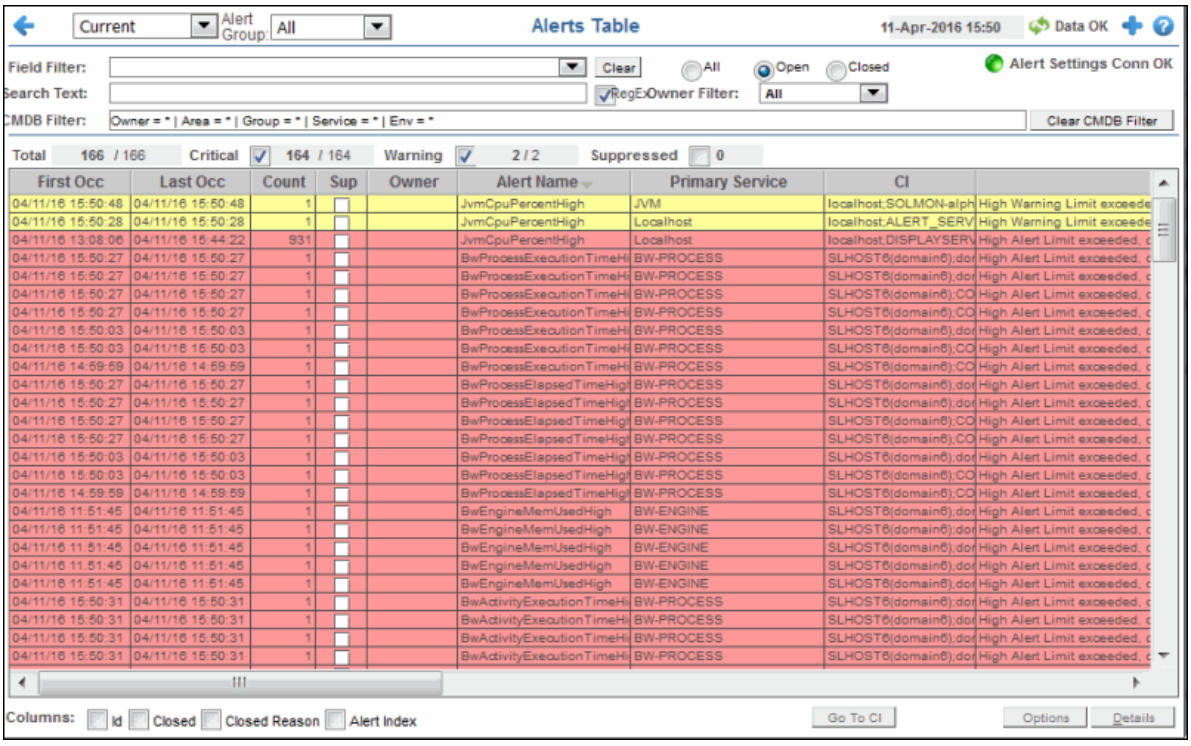

Title Bar (possible features are):

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.

The row color indicates the following:

Row Color Code:
Tables with colored rows indicate the following:

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

Fields and Data
This display includes:

Field Filter Select a table column from the drop-down menu to perform a search in: **Alert Name**, **Alert Text**, **Alert Class**, **Service**, **CI**, **Closed Reason**, **Closed**, **CompId**, **Count**, **First Occ**, **ID**, **Last Occ**, **Owner**, **Primary Service**, **Sup**, **TicketGroup**, **TicketID**.

Filters limit display content and drop-down menu selections to only those items that pass through the selected filter's criteria. If no items match the filter, you might have zero search results (an empty table).

Clear Clears the **Field Filter** and **Search Text** entries.

Search Text Enter the (case-sensitive) string to search for in the selected **Field Filter.**

CMDB Filter Shows the selected Owner, Area, Group, Service and Environment filters. By default, all components of the CMDB (*) are included in the search.

These **CMDB Filter** fields are populated when you click Open Alerts Table , which is accessible from the **Multi Area Service Views** displays, to open the **Alerts Table** in a new window. The filters selected in the **All Management Areas** and **Multi Area Service Views** displays are applied to the **Alerts Table** (that opens in the new window). NOTE: When you use the navigation tree (in the left panel) to open the **Alerts Table** display, the **Environment** filter is applied to the display if it has a value other than ***** (asterisk).

Clear CMDB Filter Clears all of the values in the **CMDB Filter** (**Owner**, **Area**, **Group**, **Service** and **Environment** filters). NOTE: This action is not applied to any other display.

RegEx Toggles the **Search Text** field to accept Regular Expressions for filtering.

All Click to show all alerts in the table: **Open** and **Closed** alerts.

Open Click to only show **Open** alerts in the table.

Closed Click to only show **Closed** alerts in the table.

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.

Alert Settings Conn OK The Alert Server connection state:
- Disconnected.
- Connected.

Total **X/Y** where **X** is the total number of alerts in the table with all selected filters applied. **Y** is the number of alerts in the table with only the **CMDB** and **Cleared** filters applied.

Critical Check to show alerts in the table that are currently in a critical state. NOTE: You must check **Critical** to see alerts that are in a critical state.

X/Y where **X** is the total number of critical alerts in the table with all selected filters applied. **Y** is the number of alerts in the table with only the **CMDB Filter** and **Cleared** filters applied.

Warning Check to show alerts in the table that are currently in a warning state. NOTE: You must check **Warning** to see alerts that are in a warning state.

X/Y where **X** is the total number of warning alerts in the table with all selected filters applied. **Y** is the number of alerts in the table with only the **CMDB** and **Cleared** filters applied.

Suppressed Check to show alerts in the table that are suppressed. The **Suppressed** count is not impacted by the **Critical** and **Warning** filters. It is impacted only by the **CMDB Filter** and the **Owner Filter**. NOTE: You must check **Suppressed** to see Suppressed alerts in the table.

Own Click to assign an Owner for the alert. This option is only visible when logged in as one of the following roles: event, full, admin, super. This option is disabled when you select a gray row. For details, see **Configure User and Role Management**.

Suppress Click to suppress the alert. This option is only visible when logged in as one of the following roles: event, full, admin, super. This option is disabled when you select a gray row. For details, see **Configure User and Role Management**.

UnSuppress Click to unsuppress the alert. This option is only visible when logged in as one of the following roles: event, full, admin, super. This option is disabled when you select a gray row or when you select a row. For details, see **Configure User and Role Management.**

Close Click to close the alert. This option is only visible to users with Administrator privileges. This option is disabled when you select a gray row or you select a row where the Primary Service is not in the $rtvManageableCompID list for the logged in user. For details, see **Configure User and Role Management**.

Alerts Table
This table lists all active alerts for the current filters. The table is empty unless you check **Critical**, **Warning**, or both. Filter the list using the search fields and drop-down menus (in the upper portion of the display). To view details about an alert, select an alert and click **Details** (in the bottom right portion of the display) to open the **Alert Detail** dialog. To view details about the CI source of the alert, select an alert and click **Go To CI** (in the bottom right portion of the display) to open its Summary display.

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

Go To CI Select an alert from the **Alerts Table**, then click **Go To CI** to view details for the selected CI in the Summary display.

Annotate Select one or more alerts from the **Alerts Table**, then click **Annotate** to open the **Set Owner and Comments** dialog and enter comments or change alert owner. This option is only visible when logged in as one of the following roles: event, full, admin, super. This option is disabled when you select a gray row or when you select a row where the Primary Service is not in the $rtvManageableCompID list for the logged in user. For details, see **Configure User and Role Management**.

ID	Lists the alert IDs, separated by semicolons, for the alerts selected from the **Alert Table**.
Source	Lists the name of the back-end Data Server reporting the alert, separated by semicolons.
Enter Owner	Enter the name of the owner for one or more alerts, click **Set Owner of One Alert** to assign the Owner, then click **Close**. By default, this field displays the current user name.
Enter Comment	Enter a comment for one or more alerts, click **Add Comment on One Alert** to apply the Comment, then click **Close**. By default, this field displays previously entered comments. The text appears in the **Comments** field for the alert.

	Set Owner	Applies the name of the alert owner in the **Enter Owner** field for one or more alerts.
	Add Comment	Applies the comment in the **Enter Comment** field for one or more alerts.
	Clear Comments	Removes all comments for one or more alerts.
	Close	Closes the dialog.
Options		Select a single alert from the **Alerts Table**, then click **Options** to open the **Alert Options** dialog. This dialog is provided for customizing your own alert options. This option is disabled when you select a gray row or more than one row.
Details		Select a single alert from the **Alerts Table**, then click **Details** to open the **Alert Detai**l window and view alert details. This option is disabled when you select a gray row or more than one row.

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"
- "Alert Administration Audit"
- "Metrics Administration"
- "RTView Cache Tables"
- "RTView Agent Admin"

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.

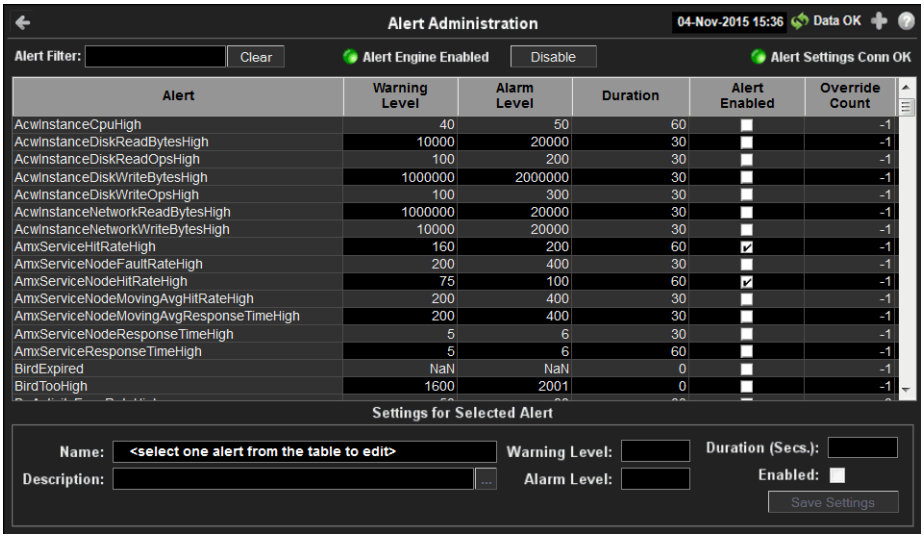

Title Bar (possible features are):

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.

Note: For more information on EMS Monitor alerts, see Appendix D, "Alert Definitions."

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**..

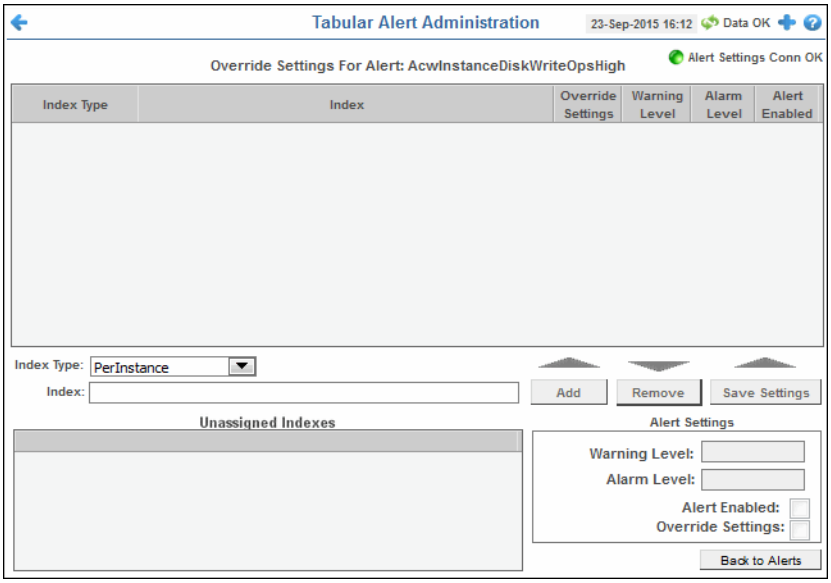

Fields and Data
This display includes:

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 a tabular alert in the **Active Alert Table** and click **Override Settings**. The **Tabular Alert Administration** display opens.

Note: Alerts that do not support overrides have a value of **-1** for the **Override Count** column and the **Override Settings** option is not present when you select such an alert.

---

2. In the **Tabular Alert Administration** display, select the Index type from the **Index Type** drop-down menu (options are populated by the type of alert you previously selected). For example, with the EMS Monitor package, select PerServerQueue, PerServerTopic or PerServer. NOTE: If you select PerServerQueue or PerServerTopic, the alert settings are applied to the queue or topic on a single server.

3. In the **Unassigned Indexes** table, select the item you want to apply an override alert setting to, click **Add** and **OK** in the confirmation dialog. After a few moments the override setting appears in the **AlertLevels** table.

4. Select the item in the **AlertLevels** table.

5. In the Alert Settings panel (lower right), if needed, modify the Warning Level and Alarm Level settings.

6. In the **Alert Settings** panel, set the following as appropriate.

- To turn on the alert for this index with the given thresholds:

  Alert Enabled Select this option.

  Override Settings Select this option.

  NOTE: To turn on an alert, both **Alert Enabled** and **Override Settings** must be selected.

- To turn off the alert for only this index (global alert thresholds will no longer apply to this index):

  Alert Enabled Deselect this option.

  Override Settings Select this option.

- To no longer evaluate this indexed alert and revert to global settings (or, optionally, Remove it if it is never to be used again):

  Alert Enabled Not used.

  Override Settings Deselect this option.

7. Click **Save Settings**. In a few moments the modifications are updated and a new record appears in the **AlertLevels** table. For example, in the following figure, the EmsServerConnectionCountHigh alert has a new override applied. New overrides increment the alert **Override Count** in the **ALERTLEVELS** table.

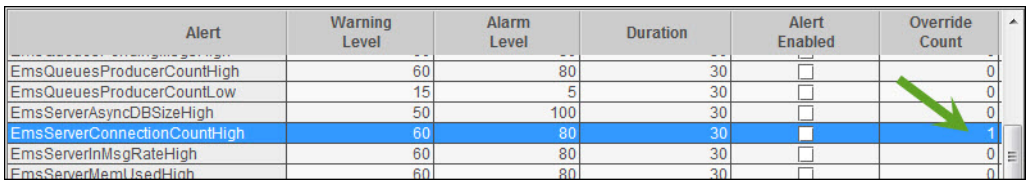

Alert	Warning Level	Alarm Level	Duration	Alert Enabled	Override Count
EmsQueuesProducerCountHigh	60	80	30	☐	0
EmsQueuesProducerCountLow	15	5	30	☐	0
EmsServerAsyncDBSizeHigh	50	100	30	☐	0
EmsServerConnectionCountHigh	60	80	30	☐	1
EmsServerInMsgRateHigh	60	80	30	☐	0
EmsServerMemUsedHigh	60	80	30	☐	0

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.

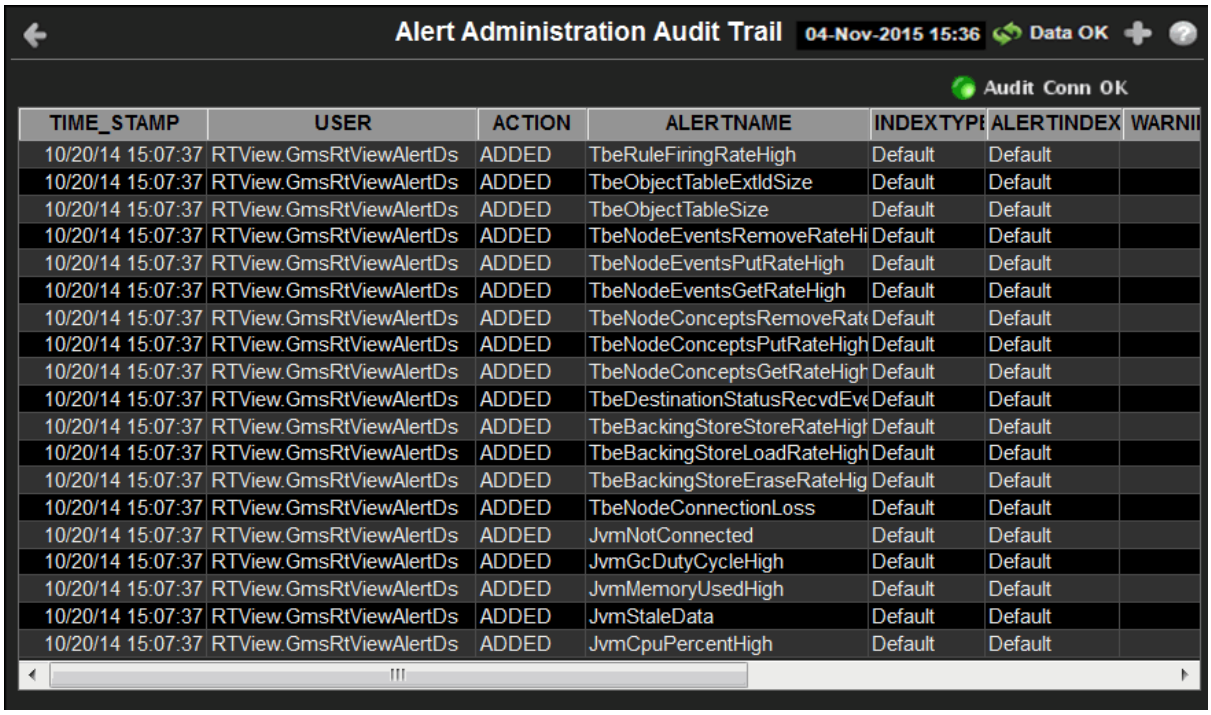

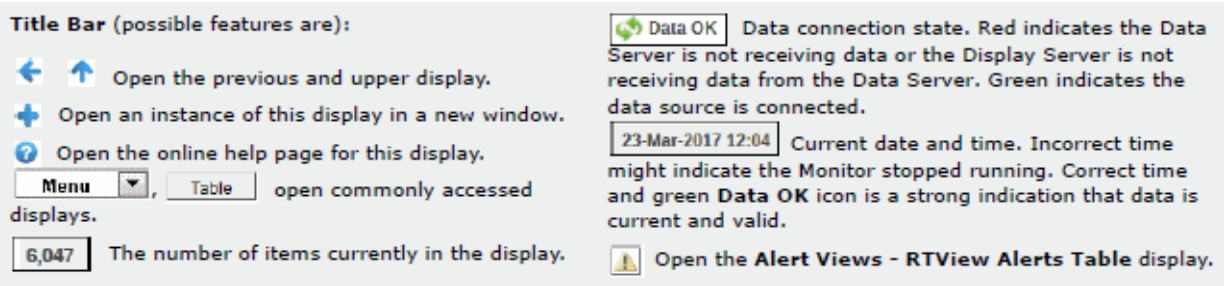

Audit Conn OK	The Alert Server connection state: Disconnected. Connected.
TIME_STAMP	The date and time of the modification.
USER	The user name of the administrator who made the modification.
ACTION	The type of modification made to the alert, such as UPDATED.
ALERTNAME	The name of the alert modified.
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 SL Technical Support.

Debugging Notes

The **Filter Objects** and **hasParent** columns were added for debugging problems related to adding and removing filtered listeners. These two columns are very specific to internal RTView structures. For example, if you make a data attachment to **Topics**, where **Name**=**"My Topic"**, an unfiltered data object would be created internally for the Topic metric, and a filtered data object would be created internally for the **Name**=**"My Topic"** row filter. The filtered data object would be setup as a child of the **Topic** metric data object. Subsequently, the **Topic** metric would have one filtered data object, and the filtered data object would have **hasParent=true**.

Also, the following JMSADM data objects (listed in the **Admin Query Key Column** and where **Last Exec Time** is **0**) are for internally created and maintained RTView tables that reside in the data source: **FTServerTable**, **ServerRouteTable**, **ServerTable** and **__admin***. These are not TIBCO metrics that are queried. Therefore, their **Last Exec Time** remains **0**, even though they are updated.

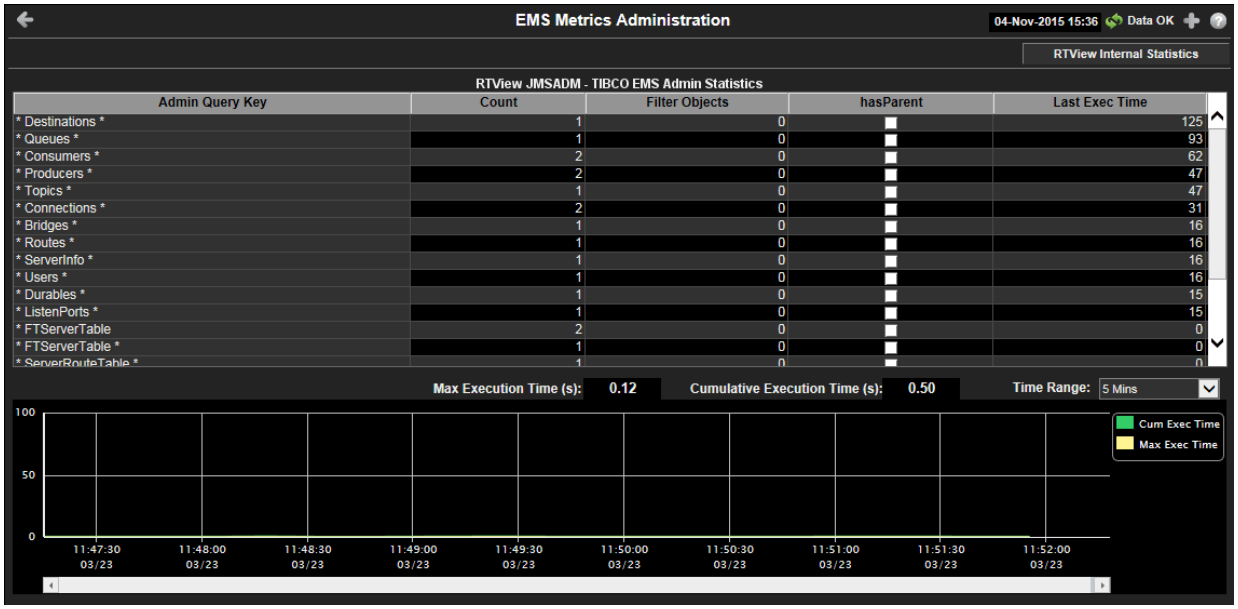

Title Bar (possible features are):

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:

RTView Internal Statistics This button opens the **RTView MBeans for Status and Timing Info** display (in a separate window), which is used primarily by SL Corporation's Technical Support team.

RTView JMSADM - TIBCO EMS Admin Statistics This table lists all JMSADM data objects. Each row in the table is a JMSADM data object. Use this data to determine the last time a TIBCO metric was queried.

- **Admin Query Key** The dsString 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.
- **Filter Objects** The number of filtered data objects in this data object.
- **hasParent** True if the data object is a filtered data object.
- **Last Exec Time** The last time a query was executed for the metric associated with this data object.

Trend Graph Traces the cumulative and maximum execution times, in seconds, for all Admin Query Keys in the table.

Cum Exec Time -- Traces the Cumulative Execution Time for all Admin Query Keys for the specified time range.

Max Exec Time -- Traces the Maximum Execution Time for all Admin Query Keys for the specified time range.

Max Execution Time The maximum execution time, in seconds, for all Admin Query Keys in the table.

Cumulative Execution Time The cumulative execution time, in seconds, for all Admin 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**. To specify a time range, click the ... button.

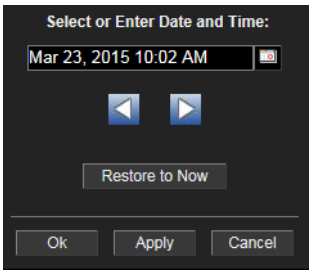

By default, the time range end point is the current time. To change the time range end point, click the button and select a date and time from the calendar or enter the date and time in the text field using the following format: **MMM dd, YYYY HH:MM**. For example, **Aug 21, 2011 12:24 PM**.

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.

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.

Click a cache table from the upper table to view cached data.

.

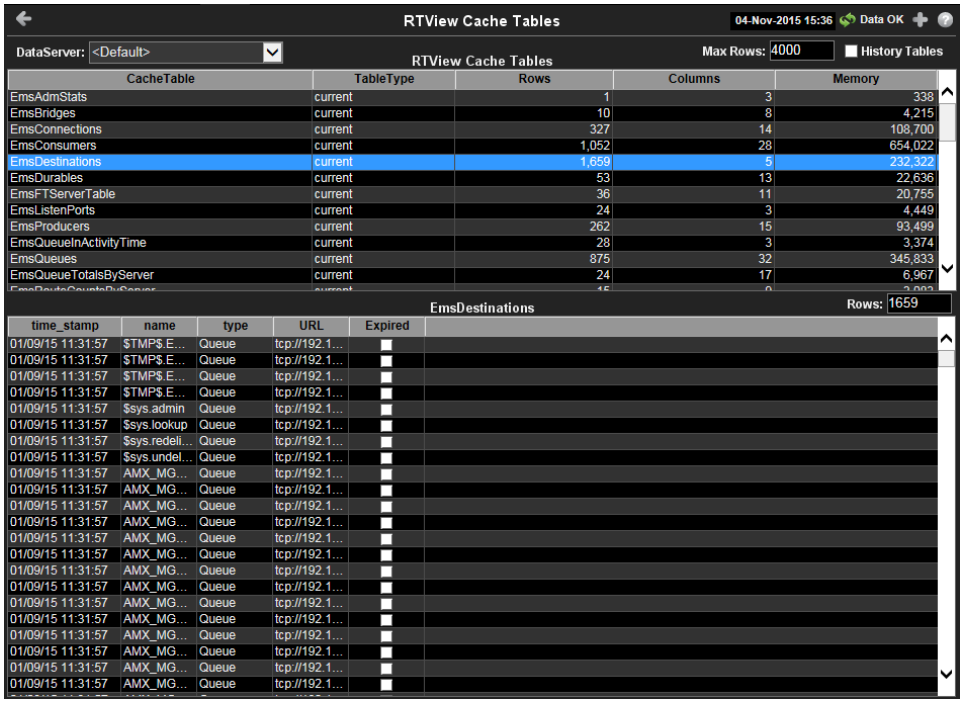

Title Bar (possible features are):

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.	

RTView Agent Admin

Verify when agent metrics were last queried by the Monitor. The data in this display is predominantly used for debugging by Technical Support.

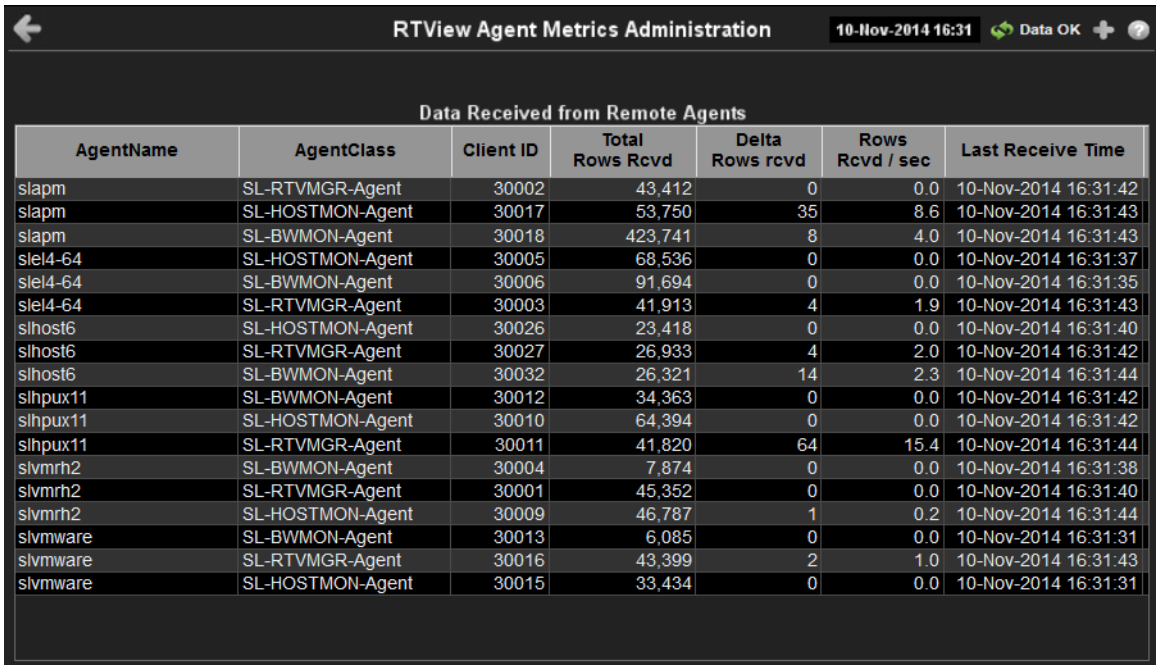

AgentName	AgentClass	Client ID	Total Rows Rcvd	Delta Rows rcvd	Rows Rcvd / sec	Last Receive Time
slapm	SL-RTVMGR-Agent	30002	43,412	0	0.0	10-Nov-2014 16:31:42
slapm	SL-HOSTMON-Agent	30017	53,750	35	8.6	10-Nov-2014 16:31:43
slapm	SL-BWMON-Agent	30018	423,741	8	4.0	10-Nov-2014 16:31:43
slel4-64	SL-HOSTMON-Agent	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

Title Bar (possible features are):

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.

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.

APPENDIX A RTView Configuration Application

This section describes settings to configure, manage and optimize your Monitor sytem using the RTView Configuration Application. This sections includes:

- "Open the RTView Configuration Application"
- "HOME Page"
- "Apply Changes"
- "Settings for RTView Central Servers"
- "Settings for Solution Package Servers"

Open the RTView Configuration Application

1. In the RTView Enterprise Monitor, click (upper right) to open the RTView Configuration Application.

Note: The icon is only visible if you are logged in as admin. You also might need to disable your browser popup blocker. If you are not logged in as admin or cannot disable your popup blocker, open the RTView Configuration Application at the following URL:

http://localhost:8068/emsample_rtvadmin

2. Login to RTView Configuration Application.

   User: **rtvadmin**

   Password: **rtvadmin**

The RTView Configuration Application **HOME** page opens.

3. Select the project you want to configure (scroll down if necessary).

HOME Page

The home page provides access to your Central Server Project and all of your Solution Package Projects. The Central Server Project allows you to configure the Enterprise Monitor Central Servers: Central Config Server, Central Alert Server, Central Alert Historian and Central Display Server. The Solution Package Projects allow you to configure all Solution Package Servers to which the Central Servers are connected.

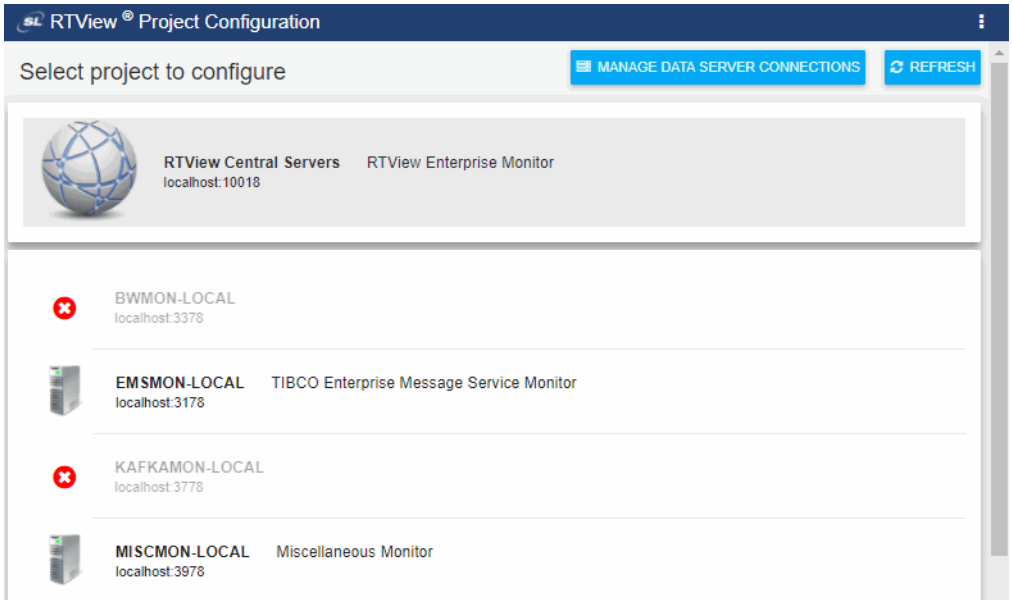

Select a project in the list to configure that project. The **MANAGE DATA SERVER CONNECTIONS** button is a shortcut to the Data Servers tab in the Central Server Project. The **REFRESH** button refreshes the Solution Package Project list.

After you have saved changes to a project, you will see the following button on the line for that project (on the home page as well as at the top of the project):

RESTART DATASERVER - This button indicates that you have an unapplied change.

Click the button to automatically restart the data server, which will require you to wait for a couple of minutes for the data server to restart. Once the data server has restarted, you can select the project to verify your changes. Note that this process only restarts the data server. Changes for the display server or historian processes will not be applied until those processes are restarted using the **RTViewEnterpriseMonitor\bin\stop_central_servers** (.bat or .sh) and **start_central_servers** (.bat or.sh) scripts and cannot be restarted from the RTView Configuration Application.

Apply Changes

Execute the **stop_central_servers** script, located in the **RTViewEnterpriseMonitor/bin** directory, then the **start_central_servers** script to restart the Central Servers Project.

Execute the **stop_data_servers** script, located in the **RTViewEnterpriseMonitor/bin** directory, then the **start_data_servers** script to restart the Solution Package Projects.

Settings for RTView Central Servers

This section describes pages and settings in the RTView Configuration Application for the RTView Central Servers. Pages are:

- "Central Server Configuration>General," next
- "Central Server Configuration>Data Servers>CONNECTIONS"
- "Central Server Configuration>Central Config Server"
- "Central Server Configuration>Central Alert Server"
- "Central Server Configuration>Central Alert Historian"
- "Central Server Configuration>Central Display Server"

The order of this section matches the RTView Configuration Application navigation tree.

TIP: If you don't see the navigation tree, click (on the left side in the title bar).

Central Server Configuration>General

The **Central Server Configuration>General** page has three tabs, **COMMON**, **CUSTOM PROPERTIES** and **ABOUT**:

- "Central Configuration Server>General>COMMON Tab," next: Use this page to connect the Central Configuration Server, the Historian and Alert Threshold databases.
- "Central Server Configuration>General>CUSTOM PROPERTIES Tab": Use this page to enter custom properties.
- "Central Server Configuration>General>ABOUT Tab": Get details about your RTView Enterprise Monitor installation from this page.

Central Configuration Server>General>COMMON Tab

Location: In the RTView Configuration Application **HOME** page, choose **RTView Central Servers**. The **General>COMMON** tab is shown by default. Use this page to connect the Central Configuration Server, the Historian and Alert Threshold databases.

The settings you make in the **COMMON** tab are applied to all Central Servers as well as all Solution Package Project Servers. For example, when you change the historian database connection in this tab, the database connection for the Central Alert Historian also changes, as well as for the Solution Package for EMS Project Server (servers\emsmon historian).

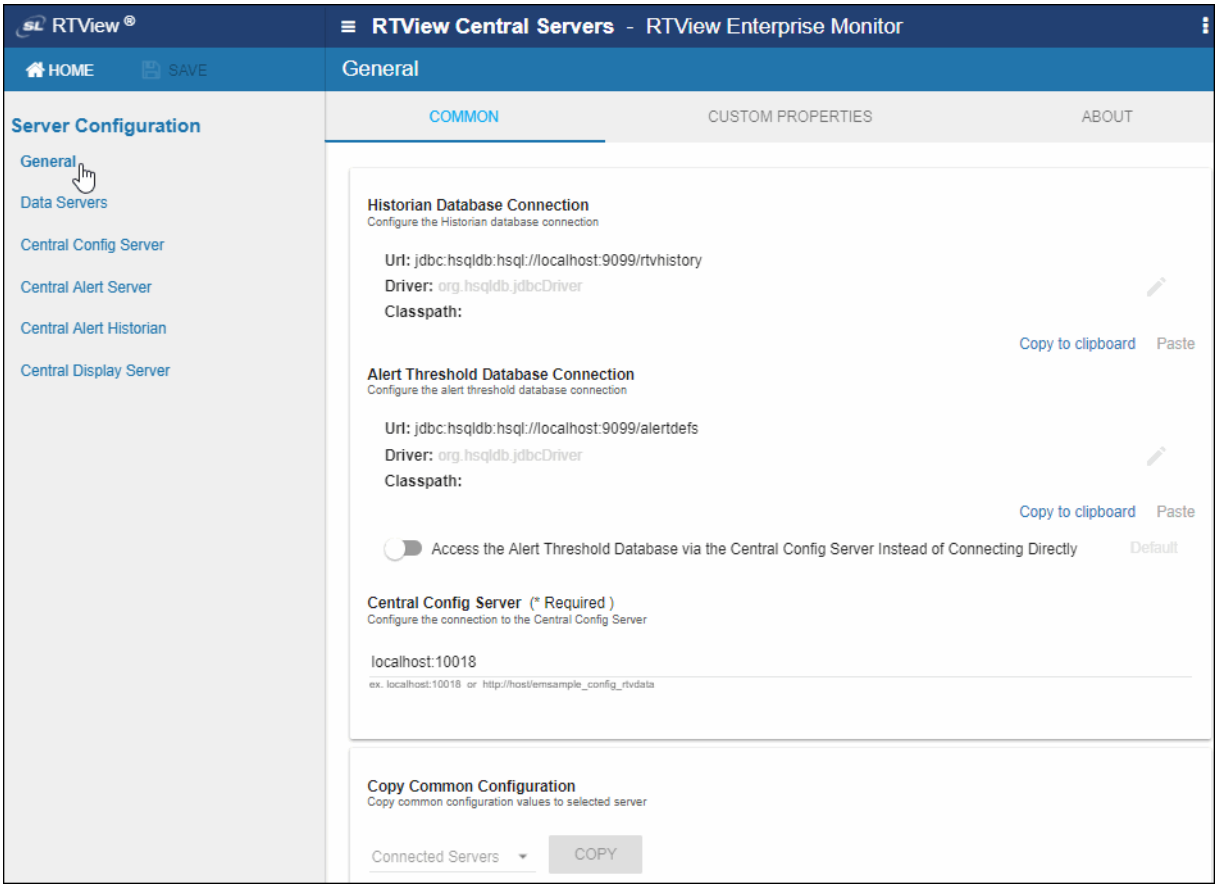

The **COMMON** tab has the following fields:

Field Name	Description
Historian Database Connection	This is the connection to use for the Historian database. See the "Configure Databases of the Central Servers" instructions on how to populate this database with the correct table schemas. **URL** - Full URL to use when connecting to this database using the specified JDBC driver. **Driver** - Fully qualified name of the driver class to use when connection to this database via JDBC. **Classpath** - The classpath to the jar containing the driver class. **Username** - (optional) User name to enter into this database when making a connection. **Password** - (optional) Password to enter into this database when making a connection.

	Run Queries Concurrently - If true, each query on the connection is run on its own execution thread. ote: This option should be used with caution since it may cause SQL errors when used with some database configurations and may degrade performance due to additional database server overhead. See your database documentation to see whether it supports concurrent queries on multiple threads. TIP: Click Copy to clipboard on any RTView Central Servers database field to copy it to the clipboard. It can be pasted on any of the other RTView Central Servers database fields.
Alert Threshold Database Connection	This is the connection to use for the Alert Threshold database. This database contains all alert settings (warning and alarm thresholds, etc). See the "Configure Databases of the Central Servers" instructions on how to populate this database with the correct table schemas. **URL** - Full URL to use when connecting to this database using the specified JDBC driver. **Driver** - Fully qualified name of the driver class to use when connection to this database via JDBC. **Classpath** - The classpath to the jar containing the driver class. **Username** - (optional) User name to enter into this database when making a connection. **Password** - (optional) Password to enter into this database when making a connection. **Run Queries Concurrently** - If true, each query on the connection is run on its own execution thread. ote: This option should be used with caution since it may cause SQL errors when used with some database configurations and may degrade performance due to additional database server overhead. See your database documentation to see whether it supports concurrent queries on multiple threads. TIP: Click **Copy** to clipboard on any RTView Central Servers database field to copy it to the clipboard. It can be pasted on any of the other RTView Central Servers database fields.
Access the Alert Threshold Database via the Central Config Server Instead of Connecting Directly.	When enabled, the Solution Package Data Servers do not connect to the Alert Threshold Database. Instead all queries go through the Central Config Server.
Central Config Server Connection	Enter the url for connecting to the Central Config Server. This field is required. Example urls: direct socket connection - localhost:10018 servlet connection - http://localhost/emsample_config_rtvdata fault tolerant pair - %PRIMARYHOST%:10018,%BACKUPHOST%:10018
Copy Common Configuration to Remote Servers	Copy the settings on this tab to a remote server. This is needed when you have a Solution Package Project that is not located under the same emsample directory as the central servers.

Central Server Configuration>General>CUSTOM PROPERTIES Tab

Use the **CUSTOM PROPERTIES** page to enter custom properties for the Central Servers.

The **CUSTOM PROPERTIES** tab has the following fields:

Field Name	Description
Custom Properties	Click + to enter a custom property. To configure a custom property, you must know the name of the associated property, the syntax for the property value and the appropriate property filter. Property values are applied in the order specified with the last value taking precedence. **Name** - the property name **Value** - the property value **Filter** - the propery filter (optional) **Comment** - a comment describing this property (optional)

Central Server Configuration>General>ABOUT Tab

The **ABOUT** tab provides details about your RTView Enterprise Monitor installation:

Field Name	Description
Location	The location of the directory where the central servers are running.
Display Name	The display name.
Version	The version of the central servers.

Central Server Configuration>Data Servers>CONNECTIONS

In the RTView Configuration Application **HOME** page, choose **RTView Central Servers>Data Servers**. Use this page to configure, enable/disable, add/remove Data Servers and also to control which solution package displays are included in the **COMPONENTS** tab of the monitor. Changes are applied after you restart the Data Server.

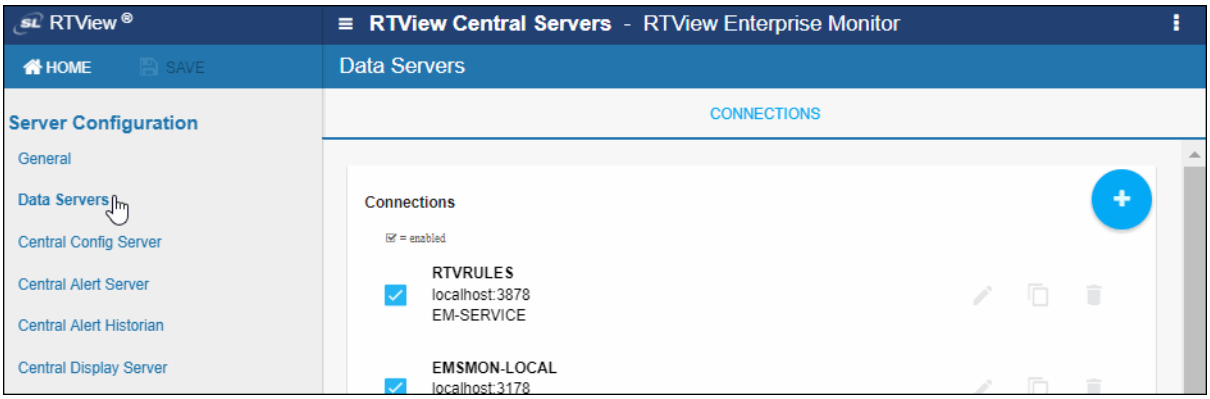

The **Central Server Configuration**>**Data Servers**>**CONNECTIONS** page has the following fields:

Field Name	Description
Connections	Add and edit connections to Solution Package Data Servers. Click to add a connection. **Name** - Connection name. This must be unique and will be the lable used on the top level of the Configuration Application for this connection. **URL** - The URL for this connection. This can be host:port or it can be a url to the rtvdata servlet. For example: **localhost:3178** **http://localhost:8068/emsmon_rtvdata** **Connection Enabled** - If true, the central servers will connect to this data server. **Select Solution Packages or CI Types hosted by this data server** - This controls which component types (CI Types) hosted by this data server will be included in the Service Model. Selecting Solution Packages includes all CI Types for that Solution Package. Optionally select **CI Types** to exclude when you do not want all CI Types to be included. **Monitor Data Server** (optional) - If true the RTView Manager will make a connection to the specified host:port so that you can monitor the process. This RTView Manager connection will use the Name. **Monitor Historian** (optional) - If true the RTView Manager will make a connection to the specified **host:port** so that you can monitor the process. This RTView Manager connection will use the Name followed by **-HISTORIAN**.
Select Mode for Including Solution Packages on the COMPONENTS Tab	Select a mode to control which Solution Package displays are included in the **COMPONENTS** tab: • **Include all solution packages in enabled connections**: This will include displays for all solution packages specified for all enabled Data Server Connections in the COMPONENTS tab of the monitor. • **Include all soluton packages in all connections**: This will include displays for all solution packages specified for all enabled and disabled Data Server Connections in the COMPONENTS tab of the monitor. • **Include all installed solution packages**: This will include displays for all installed solution packages in the COMPONENTS tab of the monitor. • **Choose solution packages to enable**: This allows you to select which solution packages will be included in the COMPONENTS tab of the monitor.

Central Server Configuration>Central Config Server

In the RTView Configuration Application **HOME** page, choose **RTView Central Servers>Central Config Serve**r. Use this page to configure the Central Config Server database connections, logging, memory and CMDB. This page has two tabs: the **CONFIGURATION SERVER** tab and the CMDB tab:

- "Central Config Server>CONFIGURATION SERVER Tab":
- "Central Config Server>CMDB Tab":

Central Config Server>CONFIGURATION SERVER Tab

Navigate to the RTView Configuration Application **HOME** page, choose **RTView Central Servers**, then choose **Central Config Server** from the navigation tree and the **CONFIGURATION SERVER** tab. Use this page to connect the Diagram Generator and Metric Explorer databases, and allocate memory.

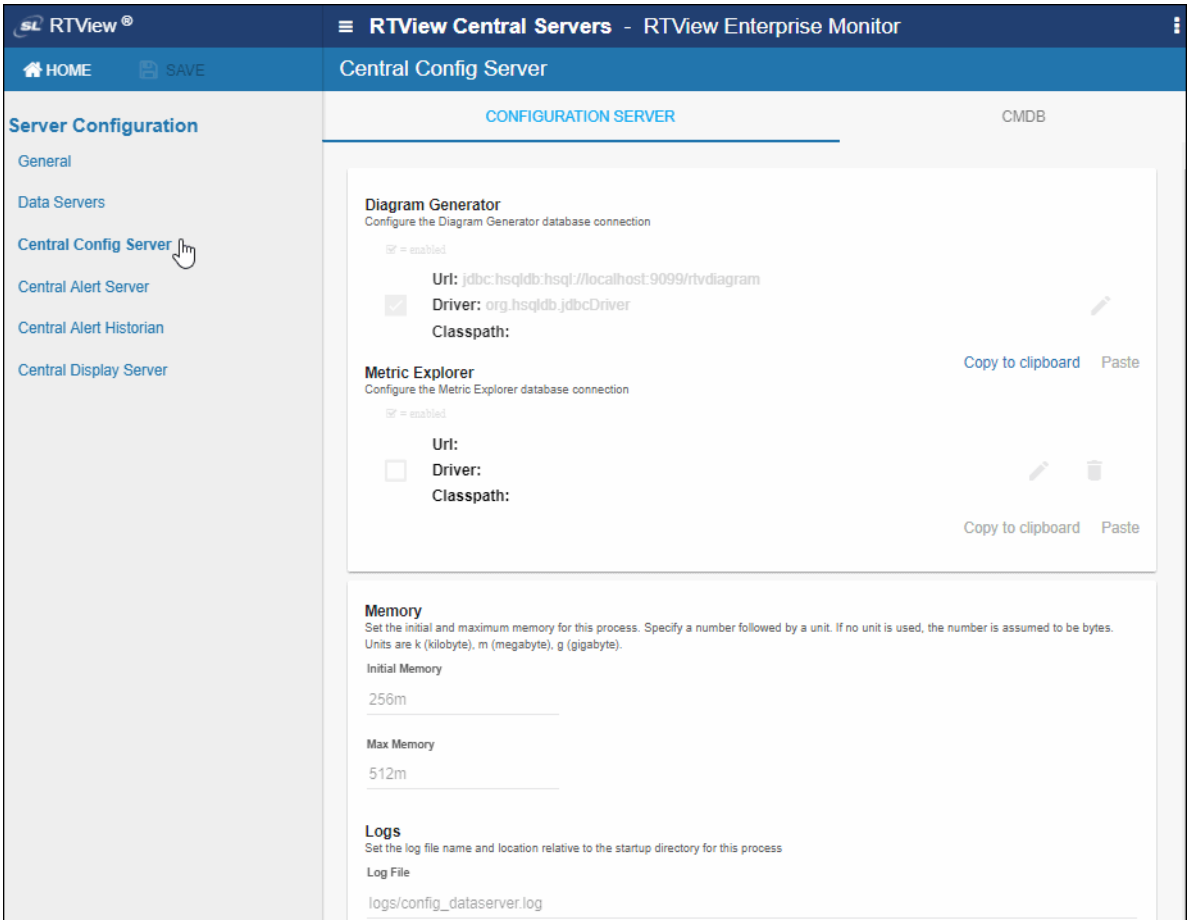

The **CONFIGURATION SERVER tab** page has the following fields:

Field Name	Description
Diagram Generator	The connection to use for the Diagram Generator. This is only needed if you will be using the Diagram Generator. **URL** - Full URL to use when connecting to this database using the specified JDBC driver. **Driver** - Fully qualified name of the driver class to use when connection to this database via JDBC. **Classpath** - The classpath to the jar containing the driver class. **Username** - (optional) User name to enter into this database when making a connection. **Password** - (optional) Password to enter into this database when making a connection.

	Run Queries Concurrently - If true, each query on the connection is run on its own execution thread. ote: This option should be used with caution since it may cause SQL errors when used with some database configurations and may degrade performance due to additional database server overhead. See your database documentation to see whether it supports concurrent queries on multiple threads. **Enabled** - Set to false to disable this database connection. TIP: Click **Copy** to clipboard on any RTView Central Servers database field to copy it to the clipboard. It can be pasted on any of the other RTView Central Servers database fields.
Metric Explorer	The connection to use for the Metric Explorer. This is only needed if you will be using the Metric Explorer. **URL** - Full URL to use when connecting to this database using the specified JDBC driver. **Driver** - Fully qualified name of the driver class to use when connection to this database via JDBC. **Classpath** - The classpath to the jar containing the driver class. **Username** - (optional) User name to enter into this database when making a connection. **Password** - (optional) Password to enter into this database when making a connection. **Run Queries Concurrently** - If true, each query on the connection is run on its own execution thread. ote: This option should be used with caution since it may cause SQL errors when used with some database configurations and may degrade performance due to additional database server overhead. See your database documentation to see whether it supports concurrent queries on multiple threads. **Enabled** - Set to false to disable this database connection. TIP: Click Copy to clipboard on any RTView Central Servers database field to copy it to the clipboard. It can be pasted on any of the other RTView Central Servers database fields.
Initial Memory*	The initial amount of memory to allocate for this process.
Max Memory*	The maximum amount of memory to allocate for this process.
Log File	The log file name and location relative to the startup directory for this process.

***Note:** Units for memory are k (kilobyte), m (megabyte), g (gigabyte). If no unit is used, the number is assumed to be bytes. Use caution when you change the memory allocation. If the memory allocation is too small the server might crash during startup (with an out of memory exception). If too large the server might eventually exceed the available CPU/memory on your system and fail.

Central Config Server>CMDB Tab

In the RTView Configuration Application **HOME** page, choose **RTView Central Servers,** then choose **Central Config Server** from the navigation tree and the **CMDB** tab. Use this page to connect and setup the CMDB (the Service Model that maps CIs being monitored).

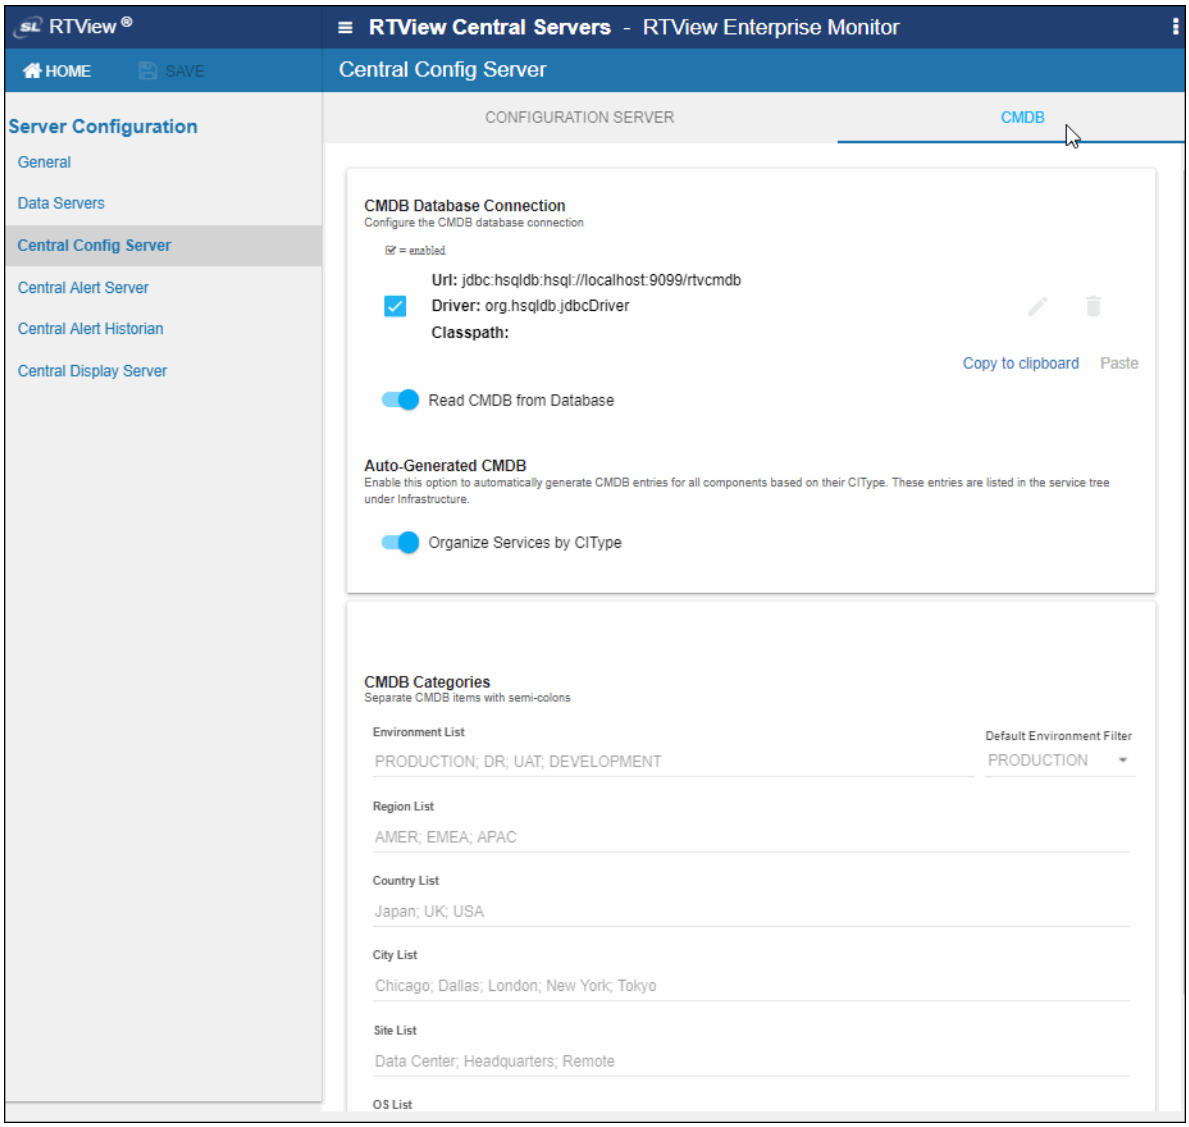

Field Name	Description
CMDB Database Connection	The database connection to use for the CMDB. This is required if you enable the Read CMDB from Database option. See "Configure Databases of the Central Servers" instructions on how to populate this database with the correct table schemas. **URL** - Full URL to use when connecting to this database using the specified JDBC driver. **Driver** - Fully qualified name of the driver class to use when connection to this database via JDBC. **Classpath** - The classpath to the jar containing the driver class.

	Username - (optional) User name to enter into this database when making a connection. **Password** - (optional) Password to enter into this database when making a connection. **Run Queries Concurrently** - If true, each query on the connection is run on its own execution thread. ote: This option should be used with caution since it may cause SQL errors when used with some database configurations and may degrade performance due to additional database server overhead. See your database documentation to see whether it supports concurrent queries on multiple threads. **Enabled** - Set to false to disable this database connection. TIP: Click Copy to clipboard on any RTView Central Servers database field to copy it to the clipboard. It can be pasted on any of the other RTView Central Servers database fields.
Read CMDB from Database	If **true**, read CMDB entries from the database defined in CMDB Database Connection
Organize Services by CIType	If **true**, enables automatic generation of CMDB entries for all components based on their CIType. When enabled, this option organizes CIs in the Service Tree based on their CI Type.
Environment List	A semi-colon delimited list of Environments to use for your CMDB entries. This populates the Environment filter list and also the list of available Environments in the CMDB Administration display.
Default Environment Filter	The initially selected value in the Environment filter field on the **SERVICE TREE** and **SERVICE VIEWS** tabs in the monitor.
Region List	A semi-colon delimited list of Regions to use for your CMDB entries. This populates the list of available Regions in the CMDB Administration display.
Country List	A semi-colon delimited list of Countries to use for your CMDB entries. This populates the list of available Countries in the CMDB Administration display.
City List	A semi-colon delimited list of Cities to use for your CMDB entries. This populates the list of available Cities in the CMDB Administration display.
Site List	A semi-colon delimited list of Site to use for your CMDB entries. This populates the list of available Site in the CMDB Administration display.
OS List	A semi-colon delimited list of Operating Systems to use for your CMDB entries. This populates the list of available Operating Systems in the CMDB Administration display.

Central Server Configuration>Central Alert Server

In the RTView Configuration Application **HOME** page, choose **RTView Central Servers** and then choose **Central Alert Server** from the navigation tree.

The **Central Alert Server** page has three tabs: the **GENERAL** tab, the **ALERTS** tab and the **DATA STORAGE** tab:

- "Central Alert Server>GENERAL Tab," next: Use this page to specify memory allocations and log files for the Central Alert Server.
- "Central Alert Server>ALERTS Tab": Use this page to configure alert notifications and persistence.
- "Central Alert Server>DATA STORAGE Tab": Use this page to configure alert history storage.

Central Alert Server>GENERAL Tab

This section describes how to allocate memory and log files for the Central Alert Server.

If you are resuming setup for a saved project: Navigate to the RTView Configuration Application **HOME** page and choose **RTView Central Servers**. The **GENERAL** tab is shown by default.

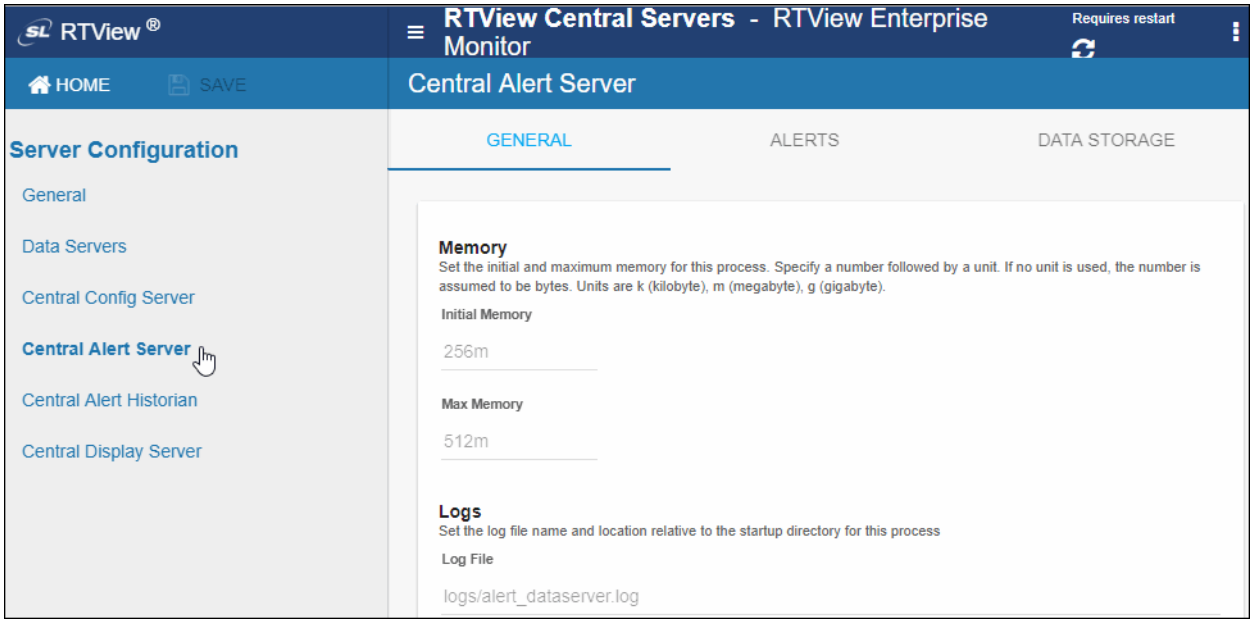

The **Central Alert Server>GENERAL tab** has the following fields:

Field Name	Description
Initial Memory*	The initial amount of memory to allocate for this process.
Max Memory*	The maximum amount of memory to allocate for this process.
Log File	The log file name and location relative to the startup directory for this process.

***Note:** Units for memory 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 (with an out of memory exception). If too large the server might eventually exceed the available CPU/ memory on your system and fail.

Central Alert Server>ALERTS Tab

This section describes how to setup alert notifications and persistence for high availability.

If you are resuming setup for a saved project: Navigate to the RTView Configuration Application **HOME** page and choose **Central Alert Server** and the **ALERTS** tab.

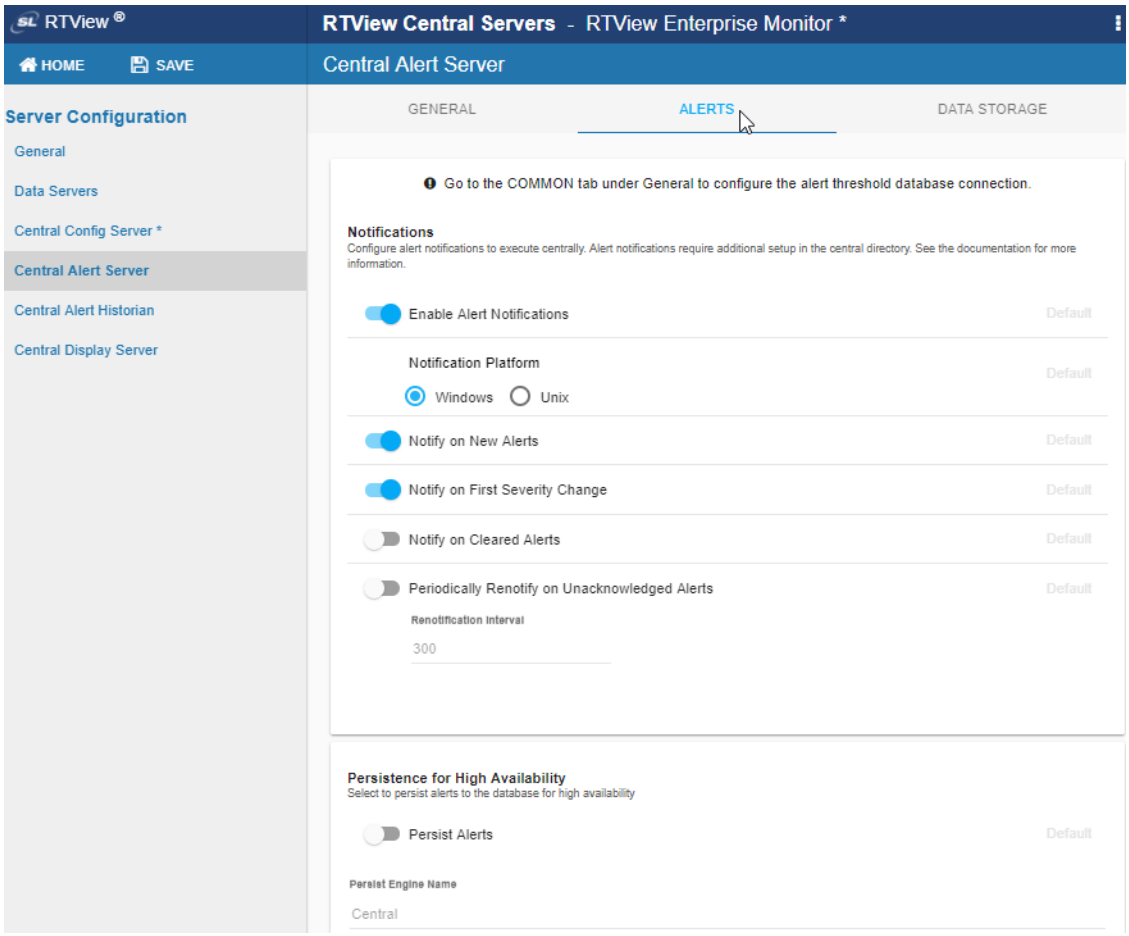

The **Central Alert Server>ALERTS tab** has the following fields:

Field Name	Description
Enable Alert Notifications	Set to true to enable alert notifications. By default, alert notifications will execute a script in the **servers/central** directory.
Notification Platform	Select the platform where the central alert server is running.
Notify on New Alerts	Set to **true** to notify on new alerts. This requires some additional setup: Copy the **my_alert_actions**(.bat or .sh) script from **RTVAPM_HOME/common/bin** to **emsample/servers/central** and modify it to execute the action you want to perform.
Notify on First Severity Change	Set to **true** to notify the first time the Severity value changes on an alert. This requires some additional setup: Copy the **my_alert_actions**(.bat or .sh) script from R**TVAPM_HOME/common/bin to emsample/servers/central** and modify it to execute the action you want to perform.
Notify on Cleared Alerts	Set to true to notify when an alert is cleared. This requires some additional setup: Copy the **my_alert_actions**(.bat or .sh) script from **RTVAPM_HOME/common/bin** to **emsample/servers/central**, rename it to **my_alert_actions.cleared**(.bat or .sh) and modify it to execute the action you want to perform.

Periodically Renotify on Unacknowledged Alerts	Set to **true** to notify on the Renotification Interval for all unacknowledged alerts. This requires some additional setup: Copy the **my_alert_actions**(.bat or .sh) script from **RTVAPM_HOME/common/bin** to **emsample/servers/central**, rename it to **my_alert_actions.renotify**(.bat or .sh) and modify it to execute the action you want to perform.
Renotification Interval	Set to the interval on which you want to renotify on unacknowledged alerts.
Persist Alerts	Set to **true** to persist the current alert table to the Alert Threshold Database. See the "Configure Databases of the Central Servers" instructions on how to populate this database with the correct table schemas.
Persist Engine Name	Assign a unique name for this data server. This is needed when multiple data servers persist alerts to the same database.

Central Alert Server>DATA STORAGE Tab

This section describes how to enable storage of alert history to the history database and create a prefix for the names of metrics stored in the history database.

If you are resuming setup for a saved project: Navigate to the RTView Configuration Application **HOME** page and choose **Central Alert Server** and the **DATA STORAGE** tab.

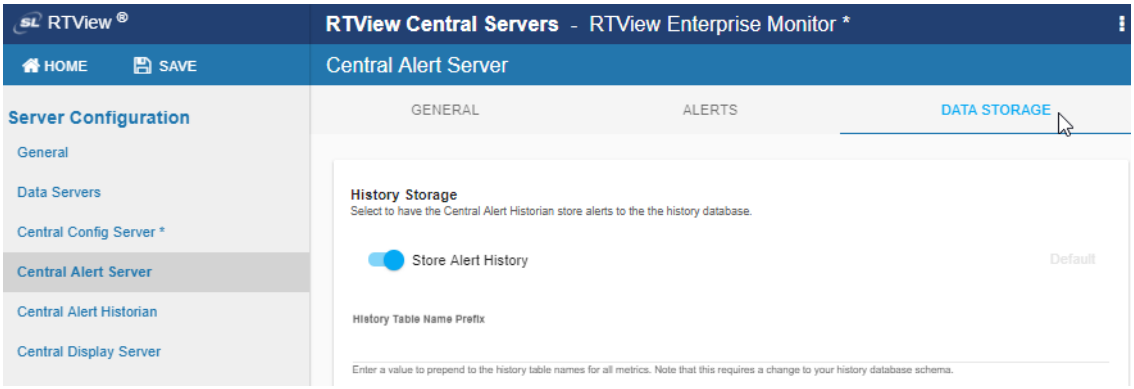

The **Central Alert Server>DATA STORAGE** tab has the following fields:

Field Name	Description
Store Alert History	Set to **true** to have the Central Alert Historian store alert history to the Historian database.
History Table Name Prefix	The **History Table Name Prefix** field allows you to define a prefix that will be added to the database table names so that EM can differentiate history data between data servers when you have multiple Central Alert Servers. 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 **RTVAPM_HOME/common/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.

Central Server Configuration>Central Alert Historian

In the RTView Configuration Application **HOME** page and choose **RTView Central Servers>Central Alert Historian**. Use this page to allocate memory and set log files for the Central Alert Historian.

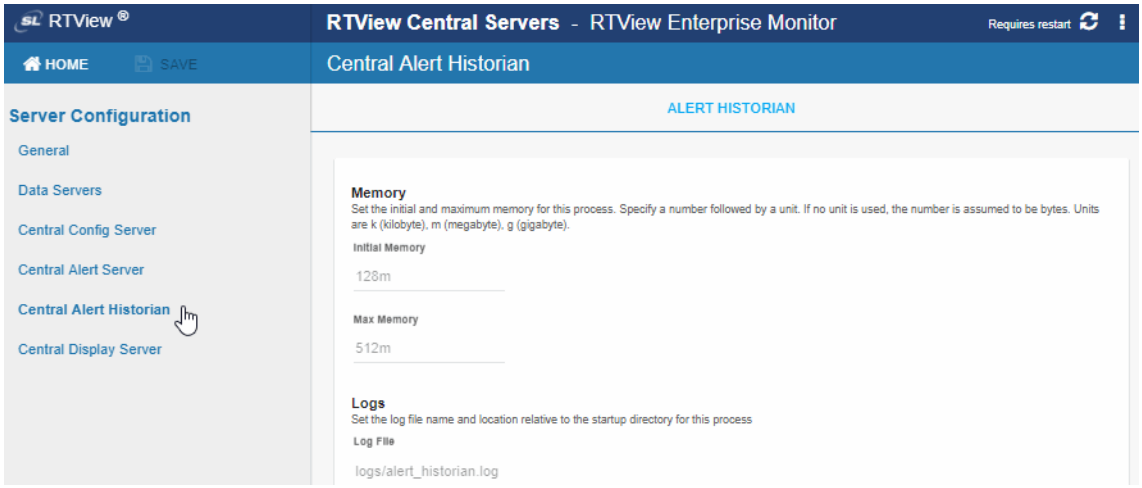

The **Central Alert Historian>ALERT HISTORIAN** tab has the following fields:

Field Name	Description
Initial Memory*	The initial amount of memory to allocate for this process.
Max Memory*	The maximum amount of memory to allocate for this process.
Log File	The log file name and location relative to the startup directory for this process.

***Note:** Units for memory 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 (with an out of memory exception). If too large the server might eventually exceed the available CPU/ memory on your system and fail.

Central Server Configuration>Central Display Server

In the RTView Configuration Application **HOME** page choose **RTView Central Servers>Central Display Server**. Use this page to choose the light or dark version of the Monitor GUI, enable the use of cross platform fonts and allocate memory and log files for the Display Server process.

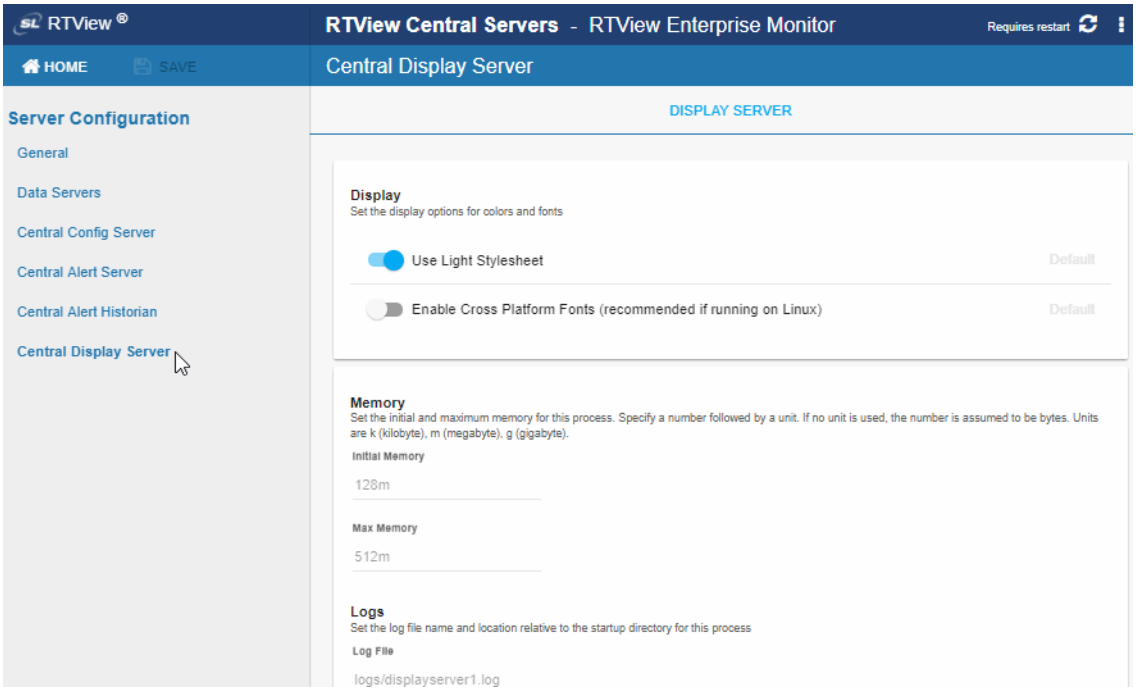

The **Central Display Server>DISPLAY SERVER** page has the following fields:

Field Name	Description
Use Light Stylesheet	Set to true to use the light colored stylesheet, false to use the dark stylesheet.
Enable Cross Platform Fonts	LINUX users might see inconsistently aligned labels in displays. To resolve, this to true.
Initial Memory*	The initial amount of memory to allocate for this process.
Max Memory*	The maximum amount of memory to allocate for this process.
Log File	The log file name and location relative to the startup directory for this process.

***Note:** Units for memory 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 (with an out of memory exception). If too large the server might eventually exceed the available CPU/ memory on your system and fail.

Settings for Solution Package Servers

This section describes the pages and settings in the RTView Configuration Application for the Solution Package Servers. Pages are:

- "Solution Package Server Configuration>General"
- "Solution Package Server Configuration>Data Server"
- "Solution Package Server Configuration>Historian"
- "Solution Package Server->Solution Package Configuration"

Solution Package Server Configuration>General

This section describes how to configure General server settings for your Solution Package projects. This page has three tabs:

- "Solution Package Server Configuration>General>GENERAL Tab": Use this page to get details about your project, set unique identifier and ports.
- "Solution Package Server Configuration>General>ALERTS Tab": Use this page to configure alert settings.
- "Solution Package Server Configuration>General>CUSTOM PROPERTIES Tab": Use this page to enter custom properties.

  Note: We use the Solution Package for RTView Manager to illustrate. Remember that each Project has it's own specifications so you might see fields or values that are not in the Solution Package for RTView Manager.

Solution Package Server Configuration>General>GENERAL Tab

In the RTView Configuration Application home page, choose **<Project Name> ->Server Configuration>General**. The **GENERAL** tab is shown by default.

Field Name	Description
URL	The Data Server URL that was used to connect to this project. This cannot be edited.
Location	The path to the Solution Package project directory. This field cannot be edited.
Version	The version of the Solution Package Data Server. This field cannot be edited.
Project Type	Displays the type of project (Standard, Sender, or ConfigClient). This field cannot be edited.
Display Name	Set the name for the project which 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.

Project ID	A default unique identifier for the project.
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.

Solution Package Server Configuration>General>ALERTS Tab

In the RTView Configuration Application home page, choose **<Project Name> ->Server Configuration>General** and go to the **ALERTS** tab.

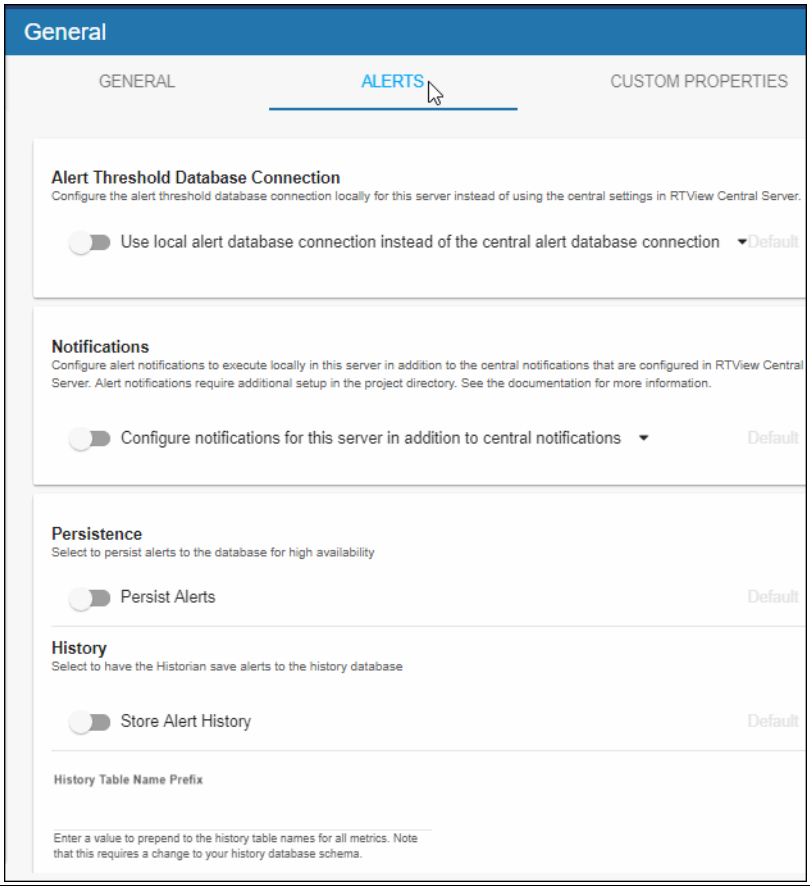

Field Name	Description
Alert Threshold Database Connection	By default, all Solution Package projects use the Alert Threshold Database connection defined under **RTView Central Servers->General->COMMON** which is the recommended setup. To use a different database for this Solution Package project, turn on the **Use local alert database connection instead of the central alert database connection** toggle and fill in the database connection information as follows:
	URL - Full URL to use when connecting to this database using the specified JDBC driver.
	Driver - Fully qualified name of the driver class to use when connection to this database via JDBC.
	Classpath - The classpath to the jar containing the driver class.

	Username - (optional) User name to enter into this database when making a connection.
	Password - (optional) Password to enter into this database when making a connection.
	Run Queries Concurrently - If true, each query on the connection is run on its own execution thread. ote: This option should be used with caution since it may cause SQL errors when used with some database configurations and may degrade performance due to additional database server overhead. See your database documentation to see whether it supports concurrent queries on multiple threads.
	Note that these limitations apply when you configure a Solution Package project to use a database other than the one specified under **RTView Central Servers->General->COMMON**:
	1. In the monitor, the **ADMIN-Alert Administration->Alert Administration** display cannot be used to set thresholds for this solution package project. Instead you must navigate to the **ADMIN-Architecture->Data Server Summary** display, select the data server for this project and click on the **Alert Admin** button.
	2. Key metrics get alert thresholds from the database defined under **RTView Central Servers->General->COMMON**.
Notifications	By default, alert notifications are disabled in the Solution Package projects since the Central Alert Server executes notifications for all alerts. To notify from this Solution Package project in addition to notifying from the Central Alert Server, turn on the Configure notifications for this server in addition to central notifications toggle and fill in the fields as follows:
	Enable Alert Notifications: Set to **true** to enable alert notifications. By default, alert notifications will execute a script in the **emsample/servers/<project directory>** directory.
	Notification Platform: Select the platform where the solution package project is running.
	Notify on New Alerts: Set to **true** to notify on new alerts. This requires some additional setup: Copy the **my_alert_actions**(.bat or .sh) script from **RTVAPM_HOME/common/bin** to **emsample/servers/<project directory>** and modify it to execute the action you want to perform.
	Notify on First Severity Change: Set to **true** to notify the first time the Severity value changes on an alert. This requires some additional setup: Copy the **my_alert_actions**(.bat or .sh) script from **RTVAPM_HOME/common/bin** to **emsample/servers/<project directory>** and modify it to execute the action you want to perform.
	Notify on Cleared Alerts: Set to **true** to notify when an alert is cleared. This requires some additional setup: Copy the **my_alert_actions**(.bat or .sh) script from **RTVAPM_HOME/common/bin** to **emsample/servers/<project directory>**, rename it to **my_alert_actions**.cleared(.bat or .sh) and modify it to execute the action you want to perform.
	Periodically Renotify on Unacknowledged Alerts: Set to **true** to notify on the Renotification Interval for all unacknowledged alerts. This requires some additional setup: Copy the **my_alert_actions**(.bat or .sh) script from **RTVAPM_HOME/common/bin** to **emsample/servers/<project directory>**, rename it to **my_alert_actions**.renotify(.bat or .sh) and modify it to execute the action you want to perform.
	Renotification Interval: Set to the interval on which you want to renotify on unacknowledged alerts.

Persist Alerts	Set to **true** to persist the current alert table to the Alert Threshold Database. See "Configure Databases of the Central Servers" instructions on how to populate this database with the correct table schemas.
History	Set to **true** to have the Historian save alert history to the history database (the Historian must be running).
History Table Name Prefix	The **History Table Name Prefix** field allows you to define a prefix that will be added to the database table names so that EM can differentiate history data between data servers when you have multiple Central Alert Servers. 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 **RTVAPM_HOME/common/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.

Solution Package Server Configuration>General>CUSTOM PROPERTIES Tab

Use the **CUSTOM PROPERTIES** page to enter custom properties for the Solution Package Servers.

The **CUSTOM PROPERTIES** tab has the following fields:

Field Name	Description
Custom Properties	Click to enter a custom property. To configure a custom property, you must know the name of the associated property, the syntax for the property value and the appropriate property filter. Property values are applied in the order specified with the last value taking precedence. **Name** - the property name. **Value** - the property value. **Filter** - the propery filter (optional). **Comment** - a comment describing this property (optional).

Solution Package Server Configuration>Data Server

This section describes how to configure data server settings for your Solution Package projects. This page has two tabs:

- "Solution Package Server Configuration>Data Server>DATA SERVER Tab": Use this page to allocate memory and set log files.
- "Solution Package Server Configuration>Data Server>SENDER Tab": Use this page to configure sender settings. This tab is only visible if the project type is sender (i.e. the project is run with the **-propfilter:sender** command line argument).

Solution Package Server Configuration>Data Server>DATA SERVER Tab

In the RTView Configuration Application home page, choose **<Project Name> ->Server Configuration>Data Server.** The **DATA SERVER** tab is shown by default.

Field Name	Description
Initial Memory*	The initial amount of memory to allocate for this process.
Max Memory*	The maximum amount of memory to allocate for this process.
Log File	The log file name and location relative to the startup directory for this process.
HTML Server Enabled	Enable the Eclipse Jetty HTML Server in the Data Server. If enabled, it will host the RTView Servlets at **http://localhost:XX70** where **XX** is the port prefix specified on the General tab. Note that you cannot disable this option if the Configuration Application is being hosted by Eclipse Jetty in the Data Server. All RTView Servlets hosted by Eclipse Jetty are automcatically configured with the correct Data Server port at runtime. The following RTView Servlets are hosted in Eclipse Jetty: rtvadmin rtvdata rtvquery rtvagent rtvpost
Use Https	Enable to use Https. This requires that you set the **Keystore File**.
	Keystore File - Set to the key store file name (including the path) that contains the certificate for your domain. This is required to use https.
	Keystore Password - Set to the password for the keystore. This field is optional.
	Key Manager Password - Set to the password for the key manager. This field is optional.

***Note:** Units for memory 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 (with an out of memory exception). If too large the server might eventually exceed the available CPU/ memory on your system and fail.

Solution Package Server Configuration>Data Server>SENDER Tab

In the RTView Configuration Application home page, choose **<Project Name> ->Server Configuration>Data Server** and go to the **SENDER** tab. This tab is only visible if the project was run with the **-propfilter:sender** command line argument.

Field Name	Description
Sender Targets	**Sender Targets**: You can specify multiple targets by adding them one at a time. All fields on the **Add Sender** dialog are required. Click the icon to open the **Add Sender** 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, http://somehost:8068/bwmon_rtvagent or http://somehost:8068/bw6mon_rtvagent).
	Targets: Select the All solution packages option.
	Enabled: Select this check box to enable the target.
Logs	The log file name and location relative to the startup directory for this process.
Sender Identifier	A unique name for the sender data server, which is typically your machine's name.

Solution Package Server Configuration>Historian

In the RTView Configuration Application home page, choose **<Project Name> ->Server Configuration>Historian**.

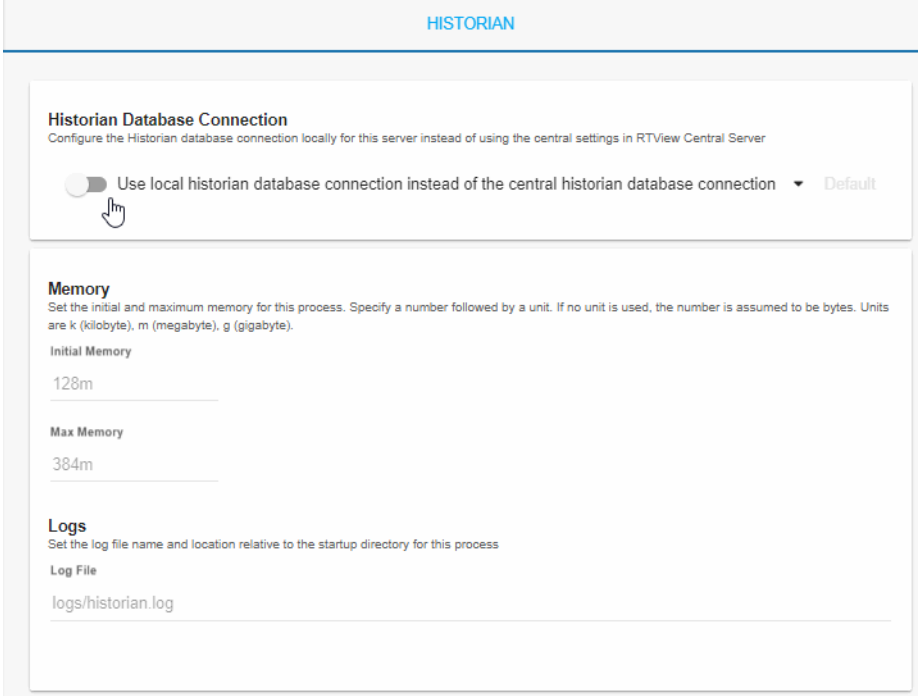

Field Name	Description
Historian Database Connection	By default,all Solution Package projects use the Historian Database connection defined under **RTView Central Servers->General->COMMON** which is the recommended setup. To use a different database for this Solution Package project, turn on the **Use local historian database connection instead of the central alert database connection** toggle and fill in the database connection information as follows:
	URL - Full URL to use when connecting to this database using the specified JDBC driver.
	Driver - Fully qualified name of the driver class to use when connection to this database via JDBC.

	Classpath - The classpath to the jar containing the driver class.
	Username - (optional) User name to enter into this database when making a connection.
	Password - (optional) Password to enter into this database when making a connection.
	Run Queries Concurrently - If **true**, each query on the connection is run on its own execution thread. **Note:** This option should be used with caution since it may cause SQL errors when used with some database configurations and may degrade performance due to additional database server overhead. See your database documentation to see whether it supports concurrent queries on multiple threads.
Initial Memory*	The initial amount of memory to allocate for this process.
Max Memory*	The maximum amount of memory to allocate for this process.
Log File	The log file name and location relative to the startup directory for this process.

***Note:** Units for memory 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 (with an out of memory exception). If too large the server might eventually exceed the available CPU/ memory on your system and fail.

Solution Package Server->Solution Package Configuration

This section describes how to add/remove and configure solution packages in your Solution Package projects. In the RTView Configuration Application home page, choose **<Project Name>**. In the naviagation tree, look for **Solution Package Configuration**. Listed under that heading are the solution packages that are included in this Solution Package project.

To add or remove a solution package, click .

Select the solution packages you want to include and click **SAVE** to close the dialog. Note that the list of solution packages in the navigation tree updates to show the solution packages you selected.

Note that you must restart your solution package project after adding a solution package before you can configure it.

APPENDIX B Monitor Scripts

This section describes scripts that are available for the Monitor as well as the **rtvservers.dat** configuration file. This section contains:

- "Scripts"
- "rtvservers.dat"

Scripts

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.

Script Name	Description
my_alert_actions.bat/sh	Sample script to define actions for alerts. Location: **rtvapm/common/bin** 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_home** This script must be executed in the directory in which it resides. Format: **rtvapm_init** (Append **.sh** on UNIX)

rtvapm_user_init.bat/sh	Initializes a user command window. Note that this script is called by **rtv_setup.bat/sh** as well as the start/stop server scripts. Location: **project directory** This script must be executed in the directory in which it resides. Format: **rtvapm_user_init** (Append **.sh** on UNIX) **Note:** This script is not available in RTViewDataCollectorSPs and RTViewDataServerSPs.
start_central_servers.bat/sh	Starts the RTView Enterprise Monitor Central 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_central_servers** (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_collector.bat/sh	Starts the RTViewDataCollectorSP server. 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_collector** (Append **.sh** on UNIX)

start_rtv.bat/sh

Starts processes in an RTView Enterprise Monitor configuration as specified in the **rtvservers.dat** configuration file.

Location:
project directory

This script must be executed in the project directory (the directory containing the **rtvservers.dat** file).

An RTView Enterprise Monitor 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 Enterprise Monitor configurations, a single RTView Enterprise Monitor configuration or a single process in an RTView Enterprise Monitor 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: 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 Enterprise Monitor configurations that are specified in the **rtvservers.dat** file.

all applies the action to all RTView Enterprise Monitor 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 Enterprise Monitor configuration specified in the **rtvservers.dat** file: **start_rtv [Configuration Name]** (Append **.sh** on UNIX) Configuration Name is the RTView Enterprise Monitor 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 Enterprise Monitor 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)
start_server.bat/sh	Starts the RTViewDataServerSP server. 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 RTView Enterprise Monitor 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 Enterprise Monitor configurations that are specified in the **rtvservers.dat** configuration file. Location: **project directory** This script must be executed in the project directory (the directory containing the **rtvservers.dat** file). This action uses defined JMX ports. An RTView Enterprise Monitor 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 Enterprise Monitor configurations, a single RTView Enterprise Monitor configuration or a single process in an RTView Enterprise Monitor 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 Enterprise Monitor 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 Enterprise Monitor configuration specified in the **rtvservers.dat** file: **status_rtv [Configuration Name]** (Append **.sh** on UNIX) Configuration Name is the RTView Enterprise Monitor 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 Enterprise Monitor 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)
stop_central_servers.bat/sh	Stops the RTView Enterprise Monitor Central 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_central_servers** (Append **.sh** on UNIX)
stop_collector.bat/sh	Stops the RTViewDataCollectorSP server. 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: **stop_collector** (Append **.sh** on UNIX)
stop_data_servers.bat/sh	Stops the data 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_data_servers** (Append **.sh** on UNIX)

stop_rtv.bat/sh	Stops processes in an RTView Enterprise Monitor configuration as specified in the **rtvservers.dat** configuration file. Location: **project directory** This script must be executed in the project directory (the directory containing the **rtvservers.dat** file). This action uses defined JMX ports. An RTView Enterprise Monitor 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 Enterprise Monitor configurations, a single RTView Enterprise Monitor configuration or a single process in an RTView Enterprise Monitor 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 Enterprise Monitor configurations that are specified in the **rtvservers.dat** file. **all** applies the action to all RTView Enterprise Monitor 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 Enterprise Monitor configuration specified in the **rtvservers.dat** file: **stop_rtv [Configuration Name]** (Append **.sh** on UNIX) Configuration Name is the RTView Enterprise Monitor 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 Enterprise Monitor 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)
stop_server.bat/sh	Stops the RTViewDataServerSP server. Location: **project directory/bin** 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 RTView Enterprise Monitor 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)
validate_install.bat/sh	Use this script if you encounter error messages when starting servers, to verify your system environment (for example, to verify that Java is installed) as well as your installation directories. Location: **<installation directory>/bin** This script must be executed in the directory in which it resides. Also, in Unix, this script checks and corrects file permissions and file formats (if, for example, the wrong unzip command was used during installation). If file permissions or formats are fixed, the script returns a count of the files fixed. Additionally, if invoked with the argument "-v" (verbose) it returns the names of the files fixed. The script returns the following information (where **<RTViewInstallation>** is your RTView installation): • In Windows Validating installation in /opt/rtview/<RTViewInstallation> ... Java installation correct. ... rtvapm installation correct. • In UNIX Validating installation in /opt/rtview/<RTViewInstallation> ... Java installation correct. ... rtvapm installation correct. ... file permissions correct. ... file formats correct.

rtvservers.dat

This section describes the **rtvservers.dat** configuration file which is used to manage your RTView Enterprise Monitor deployment and RTView Enterprise Monitor processes. This section includes:

- "Single Configuration File"
- "Multiple Configuration File"

The **rtvservers.dat** text file contains one or more RTView Enterprise Monitor configurations. An RTView Enterprise Monitor 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 Enterprise Monitor processes specified in the **rtvservers.dat** file.
- stop_rtv Stops the RTView Enterprise Monitor processes specified in the **rtvservers.dat** file.
- status_rtv Returns status information for RTView Enterprise Monitor processes specified in the **rtvservers.dat** file.

Single Configuration File

The following **rtvservers.dat** file, located in your project directory, contains a single RTView Enterprise Monitor 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 Enterprise Monitor configuration (**default** in this example).
<Project Settings Directory Location>	The RTView Enterprise Monitor 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 RTView Enterprise Monitor deployment.

For example, the following **rtvservers.dat** file, located in your project directory**/servers** directory, contains two configurations, **bwmon** and **emsmon**. Note that the project settings directory locations differ (**./bwmon** and **./emsmon**, respectively).

```
bwmon ./bwmon dataserver rundata
bwmon ./bwmon historian runhist -ds
bwmon ./bwmon 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 bwmon**. To start a single server in the **bwmon** configuration, type **start_rtv <Configuration Name> <Server Name>**. For example: **start_rtv bwmon displayserver**.

APPENDIX C Alert Definitions

This section describes alerts for Solace and their default settings.

Alert	Warning Level	Alarm Level	Duration	Enabled
SolBridgeInboundByteRateHigh The number of inbound bytes per second across the bridge has reached its maximum. Index Type: PerBridge	8000000	10000000	30	FALSE
SolBridgeInboundMsgRateHigh The number of inbound messages per second across the bridge as a whole has reached its maximum. Index Type: PerBridge	40000	50000	30	FALSE
SolBridgeOutboundByteRateHigh The number of outbound bytes per second across the bridge has reached its maximum. Index Type: PerBridge	8000000	10000000	30	FALSE
SolBridgeOutboundMsgRateHigh The number of outbound messages per second across the bridge has reached its maximum. Index Type: PerBridge	40000	50000	30	FALSE
SolClientInboundByteRateHigh The number of outbound bytes per second for the client has reached its maximum. Index Type: PerClient	8000000	10000000	30	FALSE
SolClientInboundMsgRateHigh The number of outbound messages per second for the client as a whole has reached its maximum. Index Type: PerClient	40000	50000	30	FALSE
SolClientOutboundByteRateHigh The number of outbound bytes per second for the client has reached its maximum. Index Type: PerClient	8000000	10000000	30	FALSE
SolClientOutboundMsgRateHigh The number of outbound messages per second for the client as a whole has reached its maximum. Index Type: PerClient	40000	50000	30	FALSE
SolClientSlowSubscriber One or more clients are consuming messages too slowly; endpoints may drop messages! Index Type: PerClient	1	NaN	30	FALSE

SolCspfNeighberDown State is not "OK" for one or more CSPF neighbors. Index Type: PerNeighbor	1	NaN	30	FALSE
SolEndpointPendingMsgsHigh The number of pending messages on a queue has reached its maximum. Index Type: PerEndpoint	8000	10000	30	FALSE
SolEndpointSpoolUsageHigh The endpoint is consuming too much message router memory for storing spooled messages. (Threshold units are megabytes.) Index Type: PerEndpoint	40	50	30	FALSE
SolGuaranteedMsgingHbaLinkDown For Guaranteed Messaging only, the Operational State for each HBA Fibre-Channel should be Online (e.g., not Linkdown). Index Type: PerHbaLink	0	NaN	30	FALSE
SolGuaranteedMsgingMatePortDown For Guaranteed Messaging only, the Mate Link Ports for ADB should have status OK. Index Type: PerADB	0	NaN	30	FALSE
SolGuaranteedMsgingNoMsgSpoolAdActive For Guaranteed Messaging only with Redundancy, at least one message router in an HA pair should show "AD-Active." Index Type: PerPair	0	NaN	30	FALSE
SolMsgRouterActiveDiskUtilHigh The utilization of the active disk partition for the message router is excessive. Index Type: PerAppliance	70	85	30	FALSE
SolMsgRouterByteEgressUtilHigh The egress rate (bytes/sec) utilization (current egress rate divided by max allowed) for the message router is excessive. Index Type: PerAppliance	70	85	30	FALSE
SolMsgRouterByteIngressUtilHigh The ingress rate (bytes/sec) utilization (current ingress rate divided by max allowed) for the message router is excessive. Index Type: PerAppliance	70	85	30	FALSE
SolMsgRouterConnectionUtilHigh The connection utilization for the message router (current number of connections divided by max allowed) is excessive. Index Type: PerAppliance	70	85	30	FALSE
SolMsgRouterCpuTemperatureHigh CPU temperature margin is above threshold. Index Type: PerApplianceSensor	-30	-15	30	FALSE

SolMsgRouterCspfNeighborDown Link-detect = no for CSPF neighbor. Index Type: PerAppliance	1	NaN	30	FALSE
SolMsgRouterDelvrdUnAckMsgUtilHigh The delivered unacked messages as a percentage of all messages delivered for the application is excessive. Index Type: PerAppliance	70	85	30	FALSE
SolMsgRouterFailoverDetected The backup message router in a HA pair has assumed control. Index Type: PerAppliance	1	NaN	30	FALSE
SolMsgRouterFanSensorCheckFailed The speed measured for one or more fans is below threshold. Index Type: PerApplianceSensor	5000	2657	30	FALSE
SolMsgRouterInboundByteRateHigh The number of inbound bytes per second for the message router has reached its max threshold. Index Type: PerAppliance	400000	500000	30	FALSE
SolMsgRouterInboundMsgRateHigh The number of inbound messages per second for the message router has reached its max threshold. Index Type: PerAppliance	400000	500000	30	FALSE
SolMsgRouterIngressFlowUtilHigh The ingress flow utilization (current flows divided by max allowed) for the message router is excessive. Index Type: PerAppliance	70	85	30	FALSE
SolMsgRouterInterfaceDown Link-detect = no for one or more enabled network interfaces. Index Type: PerSolInterface	NaN	NaN	30	FALSE
SolMsgRouterMsgCountUtilHigh The message count utilization for the message router is excessive. Index Type: PerAppliance	70	85	30	FALSE
SolMsgRouterMsgEgressUtilHigh The message egress rate utilization (current message egress rate divided by max allowed) for the message router is excessive. Index Type: PerAppliance	70	85	30	FALSE
SolMsgRouterMsgIngressUtilHigh The message ingress rate utilization (current message ingress rate divided by max allowed) for the message router is excessive. Index Type: PerAppliance	70	85	30	FALSE
SolMsgRouterNABUsageHigh Network Acceleration Blade memory usage is excessive. Index Type: PerNAB	60	80	30	FALSE

SolMsgRouterNotConnected The message router is not ready for collecting performance monitoring data. Index Type: PerAppliance	NaN	NaN	30	FALSE
SolMsgRouterOutboundByteRateHigh The number of outbound bytes per second for the message router has reached its max threshold. Index Type: PerAppliance	400000	500000	30	FALSE
SolMsgRouterOutboundMsgRateHigh The number of outbound messages per second for the message router has reached its max threshold. Index Type: PerAppliance	400000	500000	30	FALSE
SolMsgRouterPendingMsgsHigh The total number of pending messages for this message router has reached its maximum. Index Type: PerAppliance	400000	500000	30	FALSE
SolMsgRouterPowerSupplyFailed A power supply has failed. Index Type: PerAppliance	0	NaN	30	FALSE
SolMsgRouterSpoolUtilization The amount of spool space used for messages is excessive. Index Type: PerAppliance	70	85	30	FALSE
SolMsgRouterStandbyDiskUtilHigh The utilization of the standby disk partition for the message router is excessive. Index Type: PerAppliance	70	85	30	FALSE
SolMsgRouterSubscriptionUtilHigh The subscription utilization (current number of subscriptions divided by max allowed) for the message router is excessive. Index Type: PerAppliance	70	85	30	FALSE
SolMsgRouterSwapUsedHigh The amount of swap space used by the message router operating system is excessive. Index Type: PerAppliance	70	85	30	FALSE
SolMsgRouterSyslogAlert This alert executes when a Solace Syslog Warning or Critical message is received. To get Syslog event alerts (in RTView Enterprise Monitor or the standalone Monitor), go to the Alert Administration display and enable the **SolMsgRouterSyslog** alert.	-	-	-	-
SolMsgRouterTemperatureSensorCheckFailed A chassis temperature measurement is above threshold. Index Type: PerAppliance	40	45	30	FALSE
SolMsgRouterTranSessionCntUtilHigh The transacted session count utilization for the message router is excessive. Index Type: PerAppliance	70	85	30	FALSE

SolMsgRouterTranSessionResUtilHigh The transacted session resource utilization for the message router is excessive. Index Type: PerAppliance	70	85	30	FALSE
SolMsgRouterVoltageSensorCheckFailed A power supply voltage is high or low. Index Type: PerApplianceSesor	NaN	NaN	30	FALSE
SolVpnConnectionCountHigh The number of connections to the server has reached its maximum. Index Type: PerVPN	60	80	30	FALSE
SolVpnInboundByteRateHigh The number of inbound bytes per second for the vpn has reached its maximum. Index Type: PerVPN	8000000	10000000	30	FALSE
SolVpnInboundDiscardRateHigh The number of discarded inbound messages per second for the server is excessive. Index Type: PerVPN	1	5	30	FALSE
SolVpnInboundMsgRateHigh The number of inbound messages per second for the vpn as a whole has reached its maximum. Index Type: PerVPN	40000	50000	30	FALSE
SolVpnOutboundByteRateHigh The number of outbound bytes per second for the VPN has reached its maximum. Index Type: PerVPN	8000000	10000000	30	FALSE
SolVpnOutboundDiscardRateHigh The number of discarded outbound messages per second for the server is excessive. Index Type: PerVPN	1	5	30	FALSE
SolVpnOutboundMsgRateHigh The number of outbound messages per second for the server as a whole has reached its maximum. Index Type: PerVPN	40000	50000	30	FALSE
SolVpnPendingMsgsHigh The total number of pending messages for this destination has reached its maximum. Index Type: PerVPN	8000000	10000000	30	FALSE
SolVpnSubscriptionCountHigh The number of endpoints in this VPN has reached its maximum. Index Type: PerVPN	8000	10000	30	FALSE

APPENDIX D Third Party Notice Requirements

** Apache Tomcat is delivered for convenience only as a separate application and is licensed under the Apache License Version 2.0

** Apache HttpClient is embedded in the RTView Core libraries and is licensed under the Apache License Version 2.0

** JEval 0.9.4 is licensed under the Apache License Version 2.0

** Jetty 9.2.19 is licensed under the Apache License Version 2.0

Apache License

Version 2.0, January 2004

http://www.apache.org/licenses/

TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

1. Definitions.

"License" shall mean the terms and conditions for use, reproduction, and distribution as defined by Sections 1 through 9 of this document.

"Licensor" shall mean the copyright owner or entity authorized by the copyright owner that is granting the License.

"Legal Entity" shall mean the union of the acting entity and all other entities that control, are controlled by, or are under common control with that entity. For the purposes of this definition, "control" means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity.

"You" (or "Your") shall mean an individual or Legal Entity exercising permissions granted by this License.

"Source" form shall mean the preferred form for making modifications, including but not limited to software source code, documentation source, and configuration files.

"Object" form shall mean anyform resulting from mechanical transformation or translation of a Source form, including but not limited to compiled object code, generated documentation, and conversions to other media types.

"Work" shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in orattached to the work (an example is provided in the Appendix below)

"Derivative Works" shall mean any work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial revisions, annotations,elaborations, or other modifications represent, as a whole, an original workof authorship. For the purposes of this License, Derivative Works shallnot include works that remain separable from, or merely link (or bindby name) to the interfaces of, the Work and Derivative Works thereof.

"Contribution" shall mean anywork of authorship, including the original version of the Work and anymodifications or additions to that Work or Derivative Worksthereof, that is intentionally submitted to Licensor for inclusion inthe Work by the copyright owner or by an individual or Legal Entityauthorized to submit on behalf of the copyright owner. For the purposes of this definition, "submitted" means any form of electronic, verbal, or written communication sent to the Licensor or its representatives,including but not limited to communication on electronic mailinglists, source code control systems, and issue tracking systems that aremanaged by, or on behalf of, the Licensor for the purpose of discussingand improving the Work, but excluding communication that is conspicuously marked or otherwise designated in writing by the copyrightowner as "Not a Contribution."

"Contributor" shall meanLicensor and any individual or Legal Entity on behalf of whom a Contribution hasbeen received by Licensor and subsequently incorporated within theWork.

2. Grant of Copyright License. Subject to the terms and conditions of this License, each Contributor herebygrants to You a perpetual, worldwide, non-exclusive, no-charge,royalty-free, irrevocable copyright license to reproduce, prepareDerivative Works of, publicly display, publicly perform,sublicense, and distribute the Work and such Derivative Works in Sourceor Object form.

3. Grant of Patent License. Subject to the terms and conditions of this License, each Contributor herebygrants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable (except as stated in this section) patent license to make, have made, use, offer to sell, sell, import, and otherwise transfer the Work, where such license applies only to those patent claims licensable by such Contributor that are necessarily infringed by their Contribution(s) alone or by combinationof their Contribution(s) with the Work to which such Contribution(s) was submitted. If You institute patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent infringement,then any patent licenses granted to You under this License forthat Work shall terminate as of the date such litigation is filed.

4. Redistribution. You may reproduce anddistribute copies of the Work or Derivative Works thereof in anymedium, with or without modifications, and in Source or Objectform, provided that You meet the following conditions:

(a) You must give any other recipients of the Work or Derivative Works a copy of this License; and

(b) You must cause any modified files to carry prominent notices stating that You changed the files; and

(c) You must retain, in the Source form of any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form of the Work, excluding those notices that do not pertain to any part of the Derivative Works; and

(d) If the Work includes a "NOTICE" text file as part of its distribution, then any Derivative Works that You distribute must include a readable copy of the attribution notices contained within such NOTICE file, excluding those notices that do not pertain to any part of the Derivative Works, in at least one of the following places: within a NOTICE text file distributed as part of the Derivative Works; within the Source form or documentation, if provided along with the Derivative Works; or, within a display generated by the Derivative Works, if and wherever such third-party notices normally appear. The contents of the NOTICE file are for informational purposes only and do not modify the License. You may add Your own attribution notices within Derivative Works that You distribute, alongside or as an addendum to the NOTICE text from the Work, provided that such additional attribution notices cannot be construed as modifying the License.

You may add Your own copyright statement to Your modifications and may provide additional or different license terms and conditions for use, reproduction, or distribution of Your modifications for any such Derivative Works as a whole, provided Your use, reproduction, and distribution of the Work otherwise complies with the conditions stated in this License.

5. Submission of Contributions. Unless You explicitly state otherwise, any Contribution intentionally submitted for inclusion in the Work by You to the Licensor shall be under the terms and conditions of this License, without any additional terms or conditions. Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions.

6. Trademarks. This License does not grant permission to use the trade names, trademarks, service marks, or product names of the Licensor, except as required for reasonable and customary use in describing the origin of the Work and reproducing the content of the NOTICE file.

7. Disclaimer of Warranty. Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License.

8. Limitation of Liability. In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, including any direct, indirect, special, incidental, or consequential damages of any character arising as a result of this License or out of the use or inability to use the Work (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses), even if such Contributor has been advised of the possibility of such damages.

9. Accepting Warranty or Additional Liability. While redistributing the Work or Derivative Works thereof,You may choose to offer, and charge a fee for, acceptance ofsupport, warranty, indemnity, or other liability obligations and/orrights consistent with this License. However, in accepting suchobligations, You may act only on Your own behalf and on Your soleresponsibility, not on behalf of any other Contributor, and only ifYou agree to indemnify, defend, and hold each Contributorharmless for any liability incurred by, or claims asserted against,such Contributor by reason of your accepting any such warranty or additional liability.

END OF TERMS AND CONDITIONS

APPENDIX: How to apply the Apache License to your work.

To apply the Apache License to your work, attach the following boilerplate notice, with the fields enclosed by brackets "[]" replaced with your own identifying information. (Don't include the brackets!) The text should be enclosed in the appropriate comment syntax for the file format. We also recommend that a file or class name and description of purpose be included on the same "printed page" as the copyright notice for easier identification within third-party archives.

Copyright [yyyy] [name of copyright owner]

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at:

http://www.apache.org/licenses/LICENSE-2.0

Unless required by applicable law or agreed to in writing, software distributed under the License isdistributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANYKIND, either express or implied. See the License for the specific languagegoverning permissions and limitations under the License.

==========

** TreeMap Algorithms v1.0 is used without modifications and licensed by MPL Version 1.1. The source for TreeMap Algorithms can be obtained from http://www.cs.umd.edu/hcil/treemap/

** iTextAsian 1.0 is licensed by MPL Version 1.1 and the source can obtained from: http://itextpdf.com/download.php

MOZILLA PUBLIC LICENSE

Version 1.1

1. Definitions.

1.0.1. "Commercial Use" means distribution or otherwise making the Covered Code available to a third party.

1.1. "Contributor" means each entity that creates or contributes to the creation of Modifications.

1.2. "Contributor Version" means the combination of the Original Code, prior Modifications used by a Contributor, and the Modifications made by that particular Contributor.

1.3. "Covered Code" means the Original Code or Modifications or the combination of the Original Code andModifications, in each case including portions thereof.

1.4. "Electronic Distribution Mechanism" means a mechanism generally accepted in the software development community for the electronic transfer of data.

1.5. "Executable" means Covered Code in any form other than Source Code.

1.6. "Initial Developer" means the individual or entity identified as the Initial Developer in the SourceCode notice required by Exhibit A.

1.7. "Larger Work" means a work which combines Covered Code or portions thereof with code not governed by the terms of this License.

1.8. "License" means this document.

1.8.1. "Licensable" means having the right to grant, to the maximum extent possible, whether at the time of the initial grant or subsequently acquired, any and all of the rights conveyed herein.

1.9. "Modifications" means any addition to or deletion from the substance or structure of either the Original Code or any previous Modifications. When Covered Code is released as a series of files, a Modification is:

A. Any addition to or deletion from the contents of a file containing Original Code or previous Modifications.

B. Any new file that contains anypart of the Original Code or previous Modifications.

1.10. "Original Code" means Source Code of computer software code which is described in the Source Code notice required by Exhibit A as Original Code, and which, at the time of its release under this License is not already Covered Code governed by this License.

1.10.1. "Patent Claims" means any patent claim(s), now owned or hereafter acquired, including without limitation, method, process, and apparatus claims, in any patent Licensable by grantor.

1.11. "Source Code" means the preferred form of the Covered Code for making modifications to it, including all modules it contains, plus any associated interface definition files, scripts used to control compilation and installation of anExecutable, or source code differential comparisons against eitherthe Original Code or another well known, available Covered Code of the Contributor's choice. The Source Code can be in a compressed orarchival form, provided the appropriate decompression or de-archiving software is widely available for no charge.

1.12. "You" (or "Your") means an individual or a legal entity exercising rights under, and complyingwith all of the terms of, this License or a future version of thisLicense issued under Section 6.1. For legal entities, "You"includes any entity which controls, is controlled by, or is under common control with You. For purposes of this definition, "control"means (a) the power, direct or indirect, to cause the direction or management ofsuch entity, whether by contract or otherwise, or (b) ownershipof more than fifty percent (50%) of the outstanding shares orbeneficial ownership of such entity.

2. Source Code License.

2.1. The Initial Developer Grant.

The Initial Developer hereby grants You a world-wide, royalty-free, non-exclusive license, subject to third party intellectual property claims:

(a) under intellectual property rights (other than patent or trademark) Licensable by Initial Developer to use, reproduce, modify, display, perform, sublicense and distribute the Original Code (or portions thereof) with or without Modifications, and/or as part of a Larger Work; and

(b) under Patents Claims infringed by the making, using or selling of Original Code, to make, have made, use, practice, sell, and offer for sale, and/or otherwise dispose of the Original Code (or portions thereof).

(c) the licenses granted in this Section 2.1(a) and (b) are effective on the date Initial Developer first distributes Original Code under the terms of this License.

(d) Notwithstanding Section 2.1(b) above, no patent license is granted: 1) for code that You delete from the Original Code; 2) separate from the Original Code; or 3) for infringements caused by: i) the modification of the Original Code or ii) the combination of the Original Code with other software or devices.

2.2. Contributor Grant.

Subject to third party intellectual property claims, each Contributor hereby grants You a world-wide, royalty-free, non-exclusive license

(a) under intellectual property rights (other than patent or trademark) Licensable by Contributor, to use, reproduce, modify, display, perform, sublicense and distribute the Modifications created by such Contributor (or portions thereof) either on an unmodified basis, with other Modifications, as Covered Code and/or as part of a Larger Work; and

(b) under Patent Claims infringed by the making, using, or selling of Modifications made by that Contributor either alone and/or in combination with its Contributor Version (or portions of such combination), to make, use, sell, offer for sale, have made, and/or otherwise dispose of: 1) Modifications made by that Contributor (or portions thereof); and 2) the combination of Modifications made by that Contributor with its Contributor Version (or portions of such combination).

(c) the licenses granted in Sections2.2(a) and 2.2(b) are effective on the date Contributor first makes Commercial Use of the Covered Code.

(d) Notwithstanding Section 2.2(b) above, no patent license is granted: 1) for any code that Contributor has deleted from the Contributor Version; 2) separate from the Contributor Version;3) for infringements caused by: i) third party modifications of Contributor Version or ii) the combination of Modifications made by that Contributor with other software (except as part of the Contributor Version) or otherdevices; or 4) under Patent Claims infringed by Covered Code in theabsence of Modifications made by that Contributor.

3. Distribution Obligations.

3.1. Application of License.

The Modifications which You create or to which You contribute are governed by the terms of this License,including without limitation Section 2.2. The Source Code version of Covered Code may be distributed only under the terms of thisLicense or a future version of this License released under Section6.1, and You must include a copy of this License with every copy ofthe Source Code You distribute. You may not offer or imposeany terms on any Source Code version that alters or restricts theapplicable version of this License or the recipients' rightshereunder. However, You may include an additional document offering theadditional rights described in Section 3.5.

3.2. Availability of Source Code.

Any Modification which You create or to which You contribute must be made available in Source Code form under the terms of this License either on the same media as an Executable version or via an accepted Electronic Distribution Mechanism to anyone to whom you made an Executable version available; and if made available via Electronic Distribution Mechanism, must remain available for at least twelve (12) months after the date it initially became available, or at least six (6) months after a subsequent version of that particular Modification has been made available to such recipients. You are responsible for ensuring that the Source Code version remains available even if the Electronic Distribution Mechanism is maintained by a third party.

3.3. Description of Modifications.

You must cause all Covered Code to which You contribute to contain a file documenting the changes You made tocreate that Covered Code and the date of any change. You must includea prominent statement that the Modification is derived, directly orindirectly, from Original Code provided by the Initial Developer and including the name of the Initial Developer in (a) the Source Code,and (b) in any notice in an Executable version or related documentation in which You describe the origin or ownership of the Covered Code.

3.4. Intellectual Property Matters

(a) Third Party Claims.

If Contributor has knowledge that a license under a third party's intellectual property rights is required to exercise the rights granted by such Contributor under Sections 2.1 or 2.2, Contributor must include a text file with the Source Code distribution titled "LEGAL" which describes the claim and the party making the claim in sufficient detail that a recipient will know whom to contact. If Contributor obtains such knowledge after the Modification is made available as described in Section 3.2, Contributor shall promptly modify the LEGAL file in all copies Contributor makes available thereafter and shall take other steps (such as notifying appropriate mailing lists or newsgroups) reasonably calculated to inform those who received the Covered Code that new knowledge has been obtained.

(b) Contributor APIs.

If Contributor's Modifications include an application programming interface and Contributor has knowledge of patent licenses which are reasonably necessary to implement that API, Contributor must also include this information in the LEGAL file.

(c) Representations.

Contributor represents that, exceptas disclosed pursuant to Section 3.4(a) above, Contributorbelieves that Contributor's Modifications are Contributor'soriginal creation(s) and/or Contributor has sufficient rights togrant the rights conveyed by this License.

3.5. Required Notices.

You must duplicate the notice in Exhibit A in each file of the Source Code. If it is not possible to put such notice in a particular Source Code file due to its structure, then You must include such notice in a location (such as a relevant directory) where a user would be likely to look for such a notice. If You created one or more Modification(s) You may add your name as a Contributor to the notice described in Exhibit A. You must also duplicate this License in any documentation for the Source Code where You describe recipients' rights or ownership rights relating to Covered Code. You may choose to offer, and to charge a fee for, warranty, support, indemnity or liability obligations to one or more recipients of Covered Code. However, You may do so only on Your own behalf, and not on behalf of the Initial Developer or any Contributor. You must make it absolutely clear than any such warranty, support, indemnity or liability obligation is offered by You alone, and You hereby agree to indemnify the Initial Developer and every Contributor for any liability incurred by the Initial Developer or such Contributor as a result of warranty, support, indemnity or liability terms You offer.

3.6. Distribution of Executable Versions.

You may distribute Covered Code in Executable form only if the requirements of Section 3.1-3.5 have been met for that Covered Code, and if You include a notice stating that the Source Code version of the Covered Code is available under theterms of this License, including a description of how and whereYou have fulfilled the obligations of Section 3.2. The noticemust be conspicuously included in any notice in an Executable version,related documentation or collateral in which You describerecipients' rights relating to the Covered Code. You may distribute theExecutable version of Covered Code or ownership rights under a licenseof Your choice, which may contain terms different from this License, provided that You are in compliance with the terms of this Licenseand that the license for the Executable version does not attempt tolimit or alter the recipient's rights in the Source Code version fromthe rights set forth in this License. If You distribute the Executable version under a different license You must make it absolutely clear that any terms which differ from this License are offered by You alone, not by the Initial Developer or any Contributor. You herebyagree to indemnify the Initial Developer and every Contributorfor any liability incurred by the Initial Developer or such Contributoras a result of any such terms You offer.

3.7. Larger Works.

You may create a Larger Work by combining Covered Code with other code not governed by the terms of this License and distribute the Larger Work as a single product. In such a case, You must make sure the requirements of this License are fulfilled for the Covered Code.

4. Inability to Comply Due to Statute or Regulation.

If it is impossible for You to comply with any of the terms of this License with respect to some or all ofthe Covered Code due to statute, judicial order, or regulationthen You must: (a) comply with the terms of this License to the maximumextent possible; and (b) describe the limitations and the codethey affect. Such description must be included in the LEGAL filedescribed in Section 3.4 and must be included with all distributions of theSource Code. Except to the extent prohibited by statute orregulation, such description must be sufficiently detailed for a recipient ofordinary skill to be able to understand it.

5. Application of this License.

This License applies to code to which the Initial Developer has attached the notice in Exhibit A and to related Covered Code.

6. Versions of the License.

6.1. New Versions.

Netscape Communications Corporation ("Netscape") may publish revised and/or new versions of the License from time to time. Each version will be given a distinguishing version number.

6.2. Effect of New Versions.

Once Covered Code has been published under a particular version of the License, You may always continue to use it under the terms ofthat version. You may also choose to use suchCovered Code under the terms of any subsequent version of the Licensepublished by Netscape. No one other than Netscape has the right tomodify the terms applicable to Covered Code created under this License.

6.3. Derivative Works.

If You create or use a modified version of this License (which you may only do in order to apply it to code which is not already Covered Code governed by this License), You must (a)rename Your license so that the phrases "Mozilla","MOZILLAPL", "MOZPL", "Netscape", "MPL", "NPL" or anyconfusingly similar phrase do not appear in your license (except to note that your licensediffers from this License) and (b) otherwise make it clear that Yourversion of the license contains terms which differ from theMozilla Public License and Netscape Public License. (Filling in thename of the Initial Developer, Original Code or Contributorin the notice described in Exhibit A shall not of themselves bedeemed to be modifications of this License.)

7. DISCLAIMER OF WARRANTY.

COVERED CODE IS PROVIDED UNDER THIS LICENSE ON AN "AS IS" BASIS, WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, WARRANTIES THAT THE COVERED CODE IS FREE OF DEFECTS, MERCHANTABLE, FIT FOR A PARTICULAR PURPOSE OR NON-INFRINGING. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE COVERED CODE IS WITH YOU. SHOULD ANY COVERED CODE PROVE DEFECTIVE IN ANY RESPECT, YOU (NOT THE INITIAL DEVELOPER OR ANY OTHER CONTRIBUTOR) ASSUME THE COST OF ANY NECESSARY SERVICING, REPAIR OR CORRECTION. THIS DISCLAIMER OF WARRANTY CONSTITUTES AN ESSENTIAL PART OF THIS LICENSE. NO USE OF ANY COVERED CODE IS AUTHORIZED HEREUNDER EXCEPT UNDER THIS DISCLAIMER.

8. TERMINATION.

8.1. This License and the rights granted hereunder will terminate automatically if You fail to comply withterms herein and fail to cure such breach within 30 days of becomingaware of the breach. All sublicenses to the Covered Code which areproperly granted shall survive any termination of this License.Provisions which, by their nature, must remain in effect beyond thetermination of this License shall survive.

8.2. If You initiate litigation by asserting a patent infringement claim (excluding declatory judgment actions) against Initial Developer or a Contributor (the Initial Developeror Contributor against whom You file such action is referred to as"Participant") alleging that:

(a)such Participant's Contributor Version directly or indirectly infringes any patent, then any and allrights granted by such Participant to You under Sections 2.1and/or 2.2 of this License shall, upon 60 days notice fromParticipant terminate prospectively, unless if within 60 days after receipt ofnotice You either: (I) agree in writing to pay Participant amutually agreeable reasonable royalty for Your past and future use ofModifications made by such Participant, or (ii) withdraw Yourlitigation claim with respect to the Contributor Version against suchParticipant. If within 60 days of notice, a reasonable royalty andpayment arrangement are not mutually agreed upon in writing by theparties or the litigation claim is not withdrawn, the rights granted byParticipant to You under Sections 2.1 and/or 2.2 automaticallyterminate at the expiration of the 60 day notice period specified above.

(b)any software, hardware, or device, other than such Participant's Contributor Version, directly orindirectly infringes any patent, then any rights granted to You by suchParticipant under Sections 2.1(b) and 2.2(b) are revoked effective as ofthe date You first made, used, sold, distributed, or had made,Modifications made by that Participant.

8.3. If You assert a patent infringement claim against Participant alleging that such Participant's Contributor Version directly or indirectly infringes any patent wheresuch claim is resolved (such as by license or settlement) prior to theinitiation of patent infringement litigation, then thereasonable value of the licenses granted by such Participant underSections 2.1 or 2.2 shall be taken into account in determining the amount orvalue of any payment or license.

8.4. In the event of termination under Sections 8.1 or 8.2 above, all end user license agreements (excluding distributors and resellers) which have been validly granted by You or any distributor hereunder prior to termination shall survive termination.

9. LIMITATION OF LIABILITY.

UNDER NO CIRCUMSTANCES AND UNDER NO LEGAL THEORY, WHETHER TORT (INCLUDING NEGLIGENCE), CONTRACT, OR OTHERWISE, SHALL YOU, THE INITIAL DEVELOPER, ANY OTHER CONTRIBUTOR, OR ANY DISTRIBUTOR OF COVERED CODE, OR ANY SUPPLIER OF ANY OF SUCH PARTIES, BE LIABLE TO ANY PERSON FOR ANY INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES OF ANY CHARACTER INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF GOODWILL, WORK STOPPAGE, COMPUTER FAILURE OR MALFUNCTION, OR ANY AND ALL OTHER COMMERCIAL DAMAGES OR LOSSES, EVEN IF SUCH PARTY SHALL HAVE BEEN INFORMED OF THE POSSIBILITY OF SUCH DAMAGES. THIS LIMITATION OF LIABILITY SHALL NOT APPLY TO LIABILITY FOR DEATH OR PERSONAL INJURY RESULTING FROM SUCH PARTY'S NEGLIGENCE TO THE EXTENT APPLICABLE LAW PROHIBITS SUCH LIMITATION. SOME JURISDICTIONS DO NOT ALLOW THE EXCLUSION OR LIMITATION OF INCIDENTAL OR CONSEQUENTIAL DAMAGES, SO THIS EXCLUSION AND LIMITATION MAY NOT APPLY TO YOU.

10. U.S. GOVERNMENT END USERS.

The Covered Code is a "commercial item," as that term is defined in 48 C.F.R. 2.101 (Oct. 1995), consistingof "commercial computer software" and "commercialcomputer software documentation," as such terms are used in 48 C.F.R. 12.212 (Sept.1995). Consistent with 48 C.F.R. 12.212 and 48 C.F.R. 227.7202-1through 227.7202-4 (June 1995), all U.S. Government End Users acquireCovered Code with only those rights set forth herein.

11. MISCELLANEOUS.

This License represents the complete agreement concerning subject matter hereof. If any provision of thisLicense is held to be unenforceable, such provision shall bereformed only to the extent necessary to make it enforceable. ThisLicense shall be governed by California law provisions (except to theextent applicable law, if any, provides otherwise), excluding itsconflict-of-law provisions. With respect to disputes in which atleast one party is a citizen of, or an entity chartered or registered todo business in the United States of America, any litigationrelating to this License shall be subject to the jurisdiction of theFederal Courts of the Northern District of California, with venue lyingin Santa Clara County, California, with the losing partyresponsible for costs, including without limitation, court costs andreasonable attorneys' fees and expenses. The application of the UnitedNations Convention on Contracts for the International Sale ofGoods is expressly excluded. Any law or regulation which provides thatthe language of a contract shall be construed against the draftershall not apply to this License.

12. RESPONSIBILITY FOR CLAIMS.

As between Initial Developer and the Contributors, each party is responsible for claims and damages arising, directly or indirectly, out of its utilization of rights under this License and You agree to work with Initial Developer and Contributors to distribute such responsibility on an equitable basis. Nothing herein is intended or shall be deemed to constitute any admission of liability.

13. MULTIPLE-LICENSED CODE.

Initial Developer may designate portions of the Covered Code as "Multiple-Licensed". "Multiple-Licensed" means that the Initial Developer permits you to utilize portions of the Covered Code under Your choice of the NPL or the alternative licenses, if any, specified by the Initial Developer in the file described in Exhibit A.

EXHIBIT A -Mozilla Public License.

``The contents of this file are subject to the Mozilla Public License Version 1.1 (the "License");you may not use this file except in compliance with the License. You mayobtain a copy of the License at http://www.mozilla.org/MPL/

Software distributed under the License is distributed on an "AS IS" basis, WITHOUT WARRANTY OF ANY KIND, either express or implied. See the License for the specific language governingrights and limitations under the License.

The Original Code is _____.

The Initial Developer of the Original Code is _____.

Portions created by _____ are Copyright (C) _____ _____. All Rights Reserved.

Contributor(s): _____.

Alternatively, the contents of this file may be used under the terms of the _____ license (the "[___] License"), in which case the provisions of [_____] License are applicable instead of those above. If you wish to allow use of your version of this file only under the terms of the [_____] License and not to allow others to use your version of this file under the MPL, indicate your decision by deleting the provisions above and replace them with the notice and other provisions required by the [___] License. If you do not delete the provisions above, a recipient may use your version of this file under either the MPL or the [___] License."

[NOTE: The text of this Exhibit A may differ slightly from the text of the notices in the Source Code files of the Original Code. You should use the text of this Exhibit A rather than the text found in the Original Code Source Code for Your Modifications.]

==========

**MD Datejs

Copyright © 2006-2010 Coolite Inc.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to dealin the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE ANDNONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT,TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THESOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

==========

**JQuery

Copyright © 2009 John Resig

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE ANDNONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLEFOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OFCONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THESOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

==========

** JCalendar 1.3.2

This product uses JCalendar 1.3.2. JCalendar is distributed pursuant to the terms of the Lesser General Public License. The source code for the JCalendar may be obtained from http://www.toedter.com/en/jcalendar/index.html

==========

** BrowserLauncher2 1.3

This product uses BrowserLauncher 1.3 and is distributed pursuant to the terms of the Lesser General Public License. The source code for BrowserLauncher2 1.3 can be obtained from: http://browserlaunch2.sourceforge.net/

GNU LESSER GENERAL PUBLIC LICENSE

Version 2.1, February 1999

Copyright (C) 1991, 1999 Free Software Foundation, Inc.

59 Temple Place, Suite 330, Boston, MA 02111-1307 USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

[This is the first released version of the Lesser GPL. It also counts as the successor of the GNU Library Public License, version 2, hence the version number 2.1.]

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public Licenses are intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users.

This license, the Lesser General Public License, applies to some specially designated software packages--typically libraries--of the Free Software Foundation and other authors who decide to use it. You can use it too, but we suggest you first think carefully about whether this license or the ordinary General Public License is the better strategy to use in any particular case, based on the explanations below.

When we speak of free software, we are referring to freedom of use, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish); that you receive source code or can get it if you want it; that you can change the software and use pieces of it in new free programs; and that you are informed that you can do these things.

To protect your rights, we need to make restrictions that forbid distributors to deny you these rights or to ask you to surrender these rights. These restrictions translate to certain responsibilities for you if you distribute copies of the library or if you modify it.

For example, if you distribute copies of the library, whether gratis or for a fee, you must give the recipients all the rights that we gave you. You must make sure that they, too, receive or can get the source code. If you link other code with the library, you must provide complete object files to the recipients, so that they can relink them with the library after making changes to the library and recompiling it. And you must show them these terms so they know their rights.

We protect your rights with a two-step method: (1) we copyright the library, and (2) we offer you this license, which gives you legal permission to copy, distribute and/or modify the library.

To protect each distributor, we want to make it very clear that there is no warranty for the free library. Also, if the library is modified by someone else and passed on, the recipients should know that what they have is not the original version, so that the original author's reputation will not be affected by problems that might be introduced by others.

Finally, software patents pose a constant threat to the existence of any free program. We wish to make sure that a company cannot effectively restrict the users of a free program by obtaining a restrictive license from a patent holder. Therefore, we insist that any patent license obtained for a version of the library must be consistent with the full freedom of use specified in this license.

Most GNU software, including some libraries, is covered by the ordinary GNU General Public License. This license, the GNU Lesser General Public License, applies to certain designated libraries, and is quite different from the ordinary General Public License. We use this license for certain libraries in order to permit linking those libraries into non-free programs.

When a program is linked with a library, whether statically or using a shared library, the combination of the two is legally speaking a combined work, a derivative of the original library. The ordinary General Public License therefore permits such linking only if the entire combination fits its criteria of freedom. The Lesser General Public License permits more lax criteria for linking other code with the library.

We call this license the "Lesser" General Public License because it does Less to protect the user's freedom than the ordinary General Public License. It also provides other free software developers Less of an advantage over competing non-free programs. These disadvantages are the reason we use the ordinary General Public License for many libraries. However, the Lesser license provides advantages in certain special circumstances.

For example, on rare occasions, there may be a special need to encourage the widest possible use of a certain library, so that it becomes a de-facto standard. To achieve this, non-free programs must be allowed to use the library. A more frequent case is that a free library does the same job as widely used non-free libraries. In this case, there is little to gain by limiting the free library to free software only, so we use the Lesser General Public License.

In other cases, permission to use a particular library in non-free programs enables a greater number of people to use a large body of free software. For example, permission to use the GNU C Library in non-free programs enables many more people to use the whole GNU operating system, as well as its variant, the GNU/Linux operating system.

Although the Lesser General Public License is Less protective of the users' freedom, it does ensure that the user of a program that is linked with the Library has the freedom and the wherewithal to run that program using a modified version of the Library.

The precise terms and conditions for copying, distribution and modification follow. Pay close attention to the difference between a "work based on the library" and a "work that uses the library". The former contains code derived from the library, whereas the latter must be combined with the library in order to run.

TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License Agreement applies to any software library or other program which contains a notice placed by the copyright holder or other authorized party saying it may be distributed under the terms of this Lesser General Public License (also called "this License"). Each licensee is addressed as "you".

A "library" means a collection of software functions and/or data prepared so as to be conveniently linked with application programs (which use some of those functions and data) to form executables.

The "Library", below, refers to any such software library or work which has been distributed under these terms. A "work based on the Library" means either the Library or any derivative work under copyright law: that is to say, a work containing the Library or a portion of it, either verbatim or with modifications and/or translated straightforwardly into another language. (Hereinafter, translation is included without limitation in the term "modification".)

"Source code" for a work means the preferred form of the work for making modifications to it. For a library, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the library.

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running a program using the Library is not restricted, and output from such a program is covered only if its contents constitute a work based on the Library (independent of the use of the Library in a tool for writing it). Whether that is true depends on what the Library does and what the program that uses the Library does.

1. You may copy and distribute verbatim copies of the Library's complete source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and distribute a copy of this License along with the Library.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Library or any portion of it, thus forming a work based on the Library, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

a) The modified work must itself be a software library.

b) You must cause the files modified to carry prominent notices stating that you changed the files and the date of any change.

c) You must cause the whole of the work to be licensed at no charge to all third parties under the terms of this License.

d) If a facility in the modified Library refers to a function or a table of data to be supplied by an application program that uses the facility, other than as an argument passed when the facility is invoked, then you must make a good faith effort to ensure that, in the event an application does not supply such function or table, the facility still operates, and performs whatever part of its purpose remains meaningful.

(For example, a function in a library to compute square roots has a purpose that is entirely well-defined independent of the application. Therefore, Subsection 2d requires that any application-supplied function or table used by this function must be optional: if the application does not supply it, the square root function must still compute square roots.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Library, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Library, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Library.

In addition, mere aggregation of another work not based on the Library with the Library (or with a work based on the Library) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may opt to apply the terms of the ordinary GNU General Public License instead of this License to a given copy of the Library. To do this, you must alter all the notices that refer to this License, so that they refer to the ordinary GNU General Public License, version 2, instead of to this License. (If a newer version than version 2 of the ordinary GNU General Public License has appeared, then you can specify that version instead if you wish.) Do not make any other change in these notices.

Once this change is made in a given copy, it is irreversible for that copy, so the ordinary GNU General Public License applies to all subsequent copies and derivative works made from that copy.

This option is useful when you wish to copy part of the code of the Library into a program that is not a library.

4. You may copy and distribute the Library (or a portion or derivative of it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange.

If distribution of object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place satisfies the requirement to distribute the source code, even though third parties are not compelled to copy the source along with the object code.

5. A program that contains no derivative of any portion of the Library, but is designed to work with the Library by being compiled or linked with it, is called a "work that uses the Library". Such a work, in isolation, is not a derivative work of the Library, and therefore falls outside the scope of this License.

However, linking a "work that uses the Library" with the Library creates an executable that is a derivative of the Library (because it contains portions of the Library), rather than a "work that uses the library". The executable is therefore covered by this License. Section 6 states terms for distribution of such executables.

When a "work that uses the Library" uses material from a header file that is part of the Library, the object code for the work may be a derivative work of the Library even though the source code is not. Whether this is true is especially significant if the work can be linked without the Library, or if the work is itself a library. The threshold for this to be true is not precisely defined by law.

If such an object file uses only numerical parameters, data structure layouts and accessors, and small macros and small inline functions (ten lines or less in length), then the use of the object file is unrestricted, regardless of whether it is legally a derivative work. (Executables containing this object code plus portions of the Library will still fall under Section 6.)

Otherwise, if the work is a derivative of the Library, you may distribute the object code for the work under the terms of Section 6. Any executables containing that work also fall under Section 6, whether or not they are linked directly with the Library itself.

6. As an exception to the Sections above, you may also combine or link a "work that uses the Library" with the Library to produce a work containing portions of the Library, and distribute that work under terms of your choice, provided that the terms permit modification of the work for the customer's own use and reverse engineering for debugging such modifications.

You must give prominent notice with each copy of the work that the Library is used in it and that the Library and its use are covered by this License. You must supply a copy of this License. If the work during execution displays copyright notices, you must include the copyright notice for the Library among them, as well as a reference directing the user to the copy of this License. Also, you must do one of these things:

a) Accompany the work with the complete corresponding machine-readable source code for the Library including whatever changes were used in the work (which must be distributed under Sections 1 and 2 above); and, if the work is an executable linked with the Library, with the complete machine-readable "work that uses the Library", as object code and/or source code, so that the user can modify the Library and then relink to produce a modified executable containing the modified Library. (It is understood that the user who changes the contents of definitions files in the Library will not necessarily be able to recompile the application to use the modified definitions.)

b) Use a suitable shared library mechanism for linking with the Library. A suitable mechanism is one that (1) uses at run time a copy of the library already present on the user's computer system, rather than copying library functions into the executable, and (2) will operate properly with a modified version of the library, if the user installs one, as long as the modified version is interface-compatible with the version that the work was made with.

c) Accompany the work with a written offer, valid for at least three years, to give the same user the materials specified in Subsection 6a, above, for a charge no more than the cost of performing this distribution.

d) If distribution of the work is made by offering access to copy from a designated place, offer equivalent access to copy the above specified materials from the same place.

e) Verify that the user has already received a copy of these materials or that you have already sent this user a copy.

For an executable, the required form of the "work that uses the Library" must include any data and utility programs needed for reproducing the executable from it. However, as a special exception, the materials to be distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

It may happen that this requirement contradicts the license restrictions of other proprietary libraries that do not normally accompany the operating system. Such a contradiction means you cannot use both them and the Library together in an executable that you distribute.

7. You may place library facilities that are a work based on the Library side-by-side in a single library together with other library facilities not covered by this License, and distribute such a combined library, provided that the separate distribution of the work based on the Library and of the other library facilities is otherwise permitted, and provided that you do these two things:

a) Accompany the combined library with a copy of the same work based on the Library, uncombined with any other library facilities. This must be distributed under the terms of the Sections above.

b) Give prominent notice with the combined library of the fact that part of it is a work based on the Library, and explaining where to find the accompanying uncombined form of the same work.

8. You may not copy, modify, sublicense, link with, or distribute the Library except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, link with, or distribute the Library is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

9. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Library or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Library (or any work based on the Library), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Library or works based on it.

10. Each time you redistribute the Library (or any work based on the Library), the recipient automatically receives a license from the original licensor to copy, distribute, link with or modify the Library subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties with this License.

11. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Library at all. For example, if a patent license would not permit royalty-free redistribution of the Library by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Library.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply, and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

12. If the distribution and/or use of the Library is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Library under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

13. The Free Software Foundation may publish revised and/or new versions of the Lesser General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Library specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Library does not specify a license version number, you may choose any version ever published by the Free Software Foundation.

14. If you wish to incorporate parts of the Library into other free programs whose distribution conditions are incompatible with these, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

15. BECAUSE THE LIBRARY IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE LIBRARY, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE LIBRARY "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE LIBRARY IS WITH YOU. SHOULD THE LIBRARY PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

16. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE LIBRARY AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE LIBRARY (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE LIBRARY TO OPERATE WITH ANY OTHER SOFTWARE), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

How to Apply These Terms to Your New Libraries

If you develop a new library, and you want it to be of the greatest possible use to the public, we recommend making it free software that everyone can redistribute and change. You can do so by permitting redistribution under these terms (or, alternatively, under the terms of the ordinary General Public License).

To apply these terms, attach the following notices to the library. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

one line to give the library's name and an idea of what it does.

Copyright (C) year name of author

This library is free software; you can redistribute it and/or modify it under the terms of the GNU Lesser General Public

License as published by the Free Software Foundation; either version 2.1 of the License, or (at your option) any later version.

This library is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU Lesser General Public License for more details.

You should have received a copy of the GNU Lesser General Public License along with this library; if not, write to the Free Software Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA

Also add information on how to contact you by electronic and paper mail.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a "copyright disclaimer" for the library, if necessary. Here is a sample; alter the names:

Yoyodyne, Inc., hereby disclaims all copyright interest in the library `Frob' (a library for tweaking knobs) written by James Random Hacker.

signature of Ty Coon, 1 April 1990

Ty Coon, President of Vice

That's all there is to it!

APPENDIX E Limitations

This chapter defines the limitations experienced when using iPad Safari.

iPad Safari Limitations

- In the iPad settings for Safari, **JavaScript** must be **ON** and **Block Pop-ups** must be **OFF**. As of this writing, the Thin Client has been tested only on iOS 4.3.5 in Safari.
- The iPad does not support Adobe Flash, so the Fx graph objects (obj_fxtrend, obj_fxpie, obj_fxbar) are unavailable. The Thin Client automatically replaces the Fx graph objects with the equivalent non-Fx object (obj_trendgraph02, obj_pie, obj_bargraph). Note that the replacement objects behave the same as the Fx objects in most cases but not in all. In particular, obj_trendgraph02 does not support the sliding cursor object nor the **legendPosition** property. Custom Fx objects are not supported on the iPad.
- The Thin Client implements scrollbars for table objects and graph objects. However, unlike the scrollbars used on desktop browsers, the scrollbars used on the iPad do not have arrow buttons at each end. This can make it difficult to scroll precisely (for example, row by row) on objects with a large scrolling range.
- At full size, users may find it difficult to touch the intended display object without accidentally touching nearby objects and performing an unwanted drill-down, sort, scroll, and so forth. This is particularly true of table objects that support drill-down and also scrolling, and also in panel layouts that contain the tree navigation control. In those cases, the user may want to zoom the iPad screen before interacting with the Thin Client.
- If the iPad sleeps or auto-locks while a Thin Client display is open in Safari, or if the Safari application is minimized by clicking on the iPad's home button, the display is not updated until the iPad is awakened and Safari is reopened. In some cases it may be necessary to refresh the page from Safari's navigation bar.

Because the iPad uses a touch interface there are differences in the Thin Client appearance and behavior in iOS Safari as compared to the conventional desktop browsers that use a cursor (mouse) interface, such as Firefox and Internet Explorer. These are described below.

- Popup browser windows: An RTView object's drill-down target can be configured to open a display in a new window. In a desktop browser, when the RTView object is clicked the drill-down display is opened in a popup browser window. But in iOS Safari 4.3.5, only one page is visible at a time, so when the RTView object is touched a new page containing the drill-down display opens and fills the screen. The Safari navigation bar can be used to toggle between the currently open pages or close them.
- Mouseover text: When mouseover text and drill-down are both enabled on an RTView object (for example, a bar graph), in iOS Safari the first touch on an element in the object (for example, a bar) displays the mouseover text for that element and the second touch on the same element performs the drill-down.

- Resize Mode and Layout: By default, the Display Server runs with **resizeMode** set to **crop**. In **crop** mode, if a display is larger than the panel that contains it only a portion of the display is visible. In a desktop browser, scrollbars become available to allow the user to scroll to view the entire display. In iOS Safari, scrollbars do not appear but the display can be scrolled by dragging two fingers inside the display. (Dragging one finger scrolls the entire page, not the display).

  If the Display Server is run with **resizeMode** set to **scale** or **layout**, the display is resized to fit into the panel that contains it. If a desktop browser is resized after a display is opened, the display is resized accordingly. On the iPad, the Safari browser can only be resized by reorienting the iPad itself, between portrait mode and landscape mode.

  The panel layout feature is supported in the Thin Client. However, unlike a desktop browser which resizes to match the layout size, the size of Safari is fixed. So if the Display Server is run with **resizeMode** set to **crop** or **scale** mode, there may be unused space at the edges of the display(s) or, in **crop** mode, the panels and displays may be cropped.

  This means that **layout** mode should be used for best results on the iPad. For layout mode to be most effective, displays should use the **anchor** and **dock** object properties. Please see RTView documentation for more information.

- Scrolling: The Thin Client implements scrollbars for table objects and graph objects. The scrollbars are activated by dragging with one finger.

  If an RTView display is viewed in **crop** mode and is too large to be displayed entirely in Safari, scrollbars do not appear (as they would in a desktop browser) but the display can be scrolled by dragging with two fingers inside the display.

  Scrollbars do not ever appear in a text area control. If the text area contains more text than is visible, use the two finger drag in the text area to scroll the text.

  Regardless of the size of a listbox control, it can only display a single item (typically, the selected item). When the listbox is touched, the list of items appear in a popup list. In other words, on iOS Safari the listbox control and the combobox control behave identically.

- Context menu: The Thin Client context menu is opened by a right mouse button click in a desktop browser. It is opened in iOS Safari by touching any location on a display and holding that touch for 2 seconds. The menu appears in the top left corner of the display, regardless of where the display is touched. The items **Export Table to Excel**, **Drill Down**, and **Command** are not included on the context menu in Safari. All other items are available. The **Export Table to HTML** item is enabled if a table object is touched (unless the table object's drillDownTarget is configured to open another display). After an **Export to PDF/HTML** is performed, the exported content opens on another page in Safari. From there, the content can either be opened by another application (for example, the iBooks application opens PDF) and emailed, or it can be copied ands pasted into an email.