

Solace PubSub+ Monitor User's Guide

Version 5.3



RTView Enterprise®

Copyright © 2013-2020. Sherrill-Lubinski Corporation. All rights reserved.

RTView®

Copyright © 1998-2020. Sherrill-Lubinski Corporation. All rights reserved.

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

LIMITATIONS ON USE

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

SL, SL-GMS, GMS, RTView, 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 © 1998-2020. Sherrill-Lubinski Corporation. All rights reserved.

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

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

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



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

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

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	3
Chapter 2 - Introduction to the Monitor	7
Overview	7
Solution Package for Solace	8
Solution Package Version	8
System Requirements	8
Install PubSub+ Monitor	8
File Extraction Considerations	9
Upgrade PubSub+ Monitor	9
Solace PubSub+ Monitor v5.2.x	9
Solution Package for Solace v5.1.x	10
Solution Package for Solace v5.0 and Earlier	12
Database Table Schemas	12
RTView Configuration Application	13
Chapter 3 - Configuration	15
Setup Data Output Location	15
Output Data to PubSub+ Monitor	15
Start and Login to the Solution Package for Solace	15
Open the RTView Configuration Application	16
Configure Data Collection	17
Output Data to InfluxDB	20

Create Database and User Account	20
Output Data to Solace Broker	22
Optional Setup	22
Choose and Setup an Application Server	23
Using the Pre-configured Apache Tomcat	23
Using an Alternate Application Server	25
Modify Default Polling Rates for Solace Caches	26
Modify Default Settings for Storing Historical Data	26
Define the Storage of In Memory History	27
Define Compaction Rules	27
Define Duration	28
Enable/Disable Storage of Historical Data	29
Define Prefix for All History Table Names	29
Change Port Assignments	30
Configure Alert & Historical Database Connections	31
Troubleshoot	34
Log Files for Solace	35
JAVA_HOME	35
Permissions	35
Network/DNS	35
Data Not Received from Data Server	35
Obtain SEMP Schemas	36
 Chapter 4 - Additional Configurations	 37
Solace Event Module	37
Introduction	37
Configure PubSub+ Message Broker & Syslog Destination	37
Configure Solace Event Module	38
Solace Event Module Caches and Alerts	39
The SolEventModuleAlerts Cache	40
Solace Event Module Logging	41
High Availability	41
HA Architecture	41
Data Server HA	41
Display Server HA (Classic UI-RTView Manager Only)	42
HTML UI HA (Solace PubSub+ Monitor UI)	42
Historian HA	42
Requirements	42
Configure HA	42
Verify HA Setup	43
Primary Data Server Log File	44
Backup Data Server Log File	44
Primary Historian Log File	44
Backup Historian Log File	44
Primary Display Server Log File	45
Backup Display Server Log File	45

Property Editor REST API	45
Import Initial Properties & Connections into Configuration Application	46
Automate Connection Updates	47
Encrypt Property Text	48
Design Notes.....	49
Supported API Actions.....	49
Filenames.....	49
Sample json	49
Adding, Editing, Deleting JsonPrimitive Properties	49
Adding and Editing JsonObject Properties.....	50
Deleting JsonObject Properties	50
Updating vs. Restarting Data Servers	51
High Availability	51
Create Custom Alerts	51
Defining Alerts.....	52
Deploying Alerts	55
Verifying Alerts	55
Chapter 5 - Configure Alert Notification	57
Run a Script	59
Execute Java Code	59
Customizing the Custom Command Handler	60
Add Email Notification	60
Send SNMP Trap	61
Run Command String	61
Conditional Filter.....	62
Chapter 6 - Using the Monitor	63
Login to Solution Package for Solace	63
User Permissions	63
Overview	65
Graphic Elements.....	65
Heatmaps.....	65
Tables.....	67
Trend Graphs.....	71
Icons and Buttons	72
Displays	73
Brokers	73
Brokers Overview	73
Brokers Heatmap	75
Brokers Table	77
Broker Summary	84
Broker Sensors	86
Broker Provisioning.....	87
Broker Interface.....	89

Brokers Message Spool	91
CSPF Neighbors	93
Neighbors Table	93
CSPF Neighbors Diagram	95
Neighbors Summary	97
VPNs	99
VPNs Heatmap	99
VPNs Table	103
VPNs Summary	106
Clients	109
Clients Table	109
Client Summary	114
Bridges	117
Bridges Table	117
Bridges Diagram	120
Bridge Summary	121
Endpoints	124
Endpoints Table	124
Endpoint Summary	126
Capacity	129
Broker Capacity Table	129
Broker Capacity - Summary	131
Broker Capacity Trends	132
Syslog Events	133
Syslog Events Table	134
Syslog Event Summary	135
.....	138
Drill Down Displays	138
Alerts History Table - HTML	138
Alerts Table by Component - HTML	139
Alert Detail for Component - HTML	140
Alert Configuration for Component - HTML	142
Alerts	143
Alerts Table	144
Admin	145
Alert Administration	145
Alert Overrides Admin	147
Cache Table	148
Chapter 7 - RTView Manager	151
Login to RTView Manager	151
Displays	153
Tomcat Displays	153
Tomcat Overview	154
Tomcat Servers Heatmap	155
Single Tomcat Server	155

.....	156
Tomcat Applications Heatmap	157
Single Tomcat App.....	158
JVM Processes Displays.....	160
JVM Overview	160
JVMs Table	161
JVMs Heatmap	162
JVM Summary.....	164
JVM System Properties.....	165
JVM GC Trends.....	166
RTView Servers Displays	168
Data Servers	168
Data Server Summary	169
Historian Servers.....	169
‘Drilldowns’ Displays.....	170
Alerts History Table	170
Alerts Table by Component	171
Alert Detail for Component	173
Alerts Displays	174
Alerts Table	175
Admin Displays.....	176
Alert Administration	176
Alert Overrides Admin	178
Cache Table.....	180
Modify RTView Manager Settings.....	181
Open the RTView Configuration Application for RTView Manager	182
Modify Connections for Data Collection	183
Modify Default Polling Rates for RTView Manager Caches	185
Modify Default Settings for Storing Historical Data.....	185
Define the Storage of In Memory History	186
Define Compaction Rules	187
Define Duration.....	188
Enable/Disable Storage of Historical Data	189
Define Prefix for All History Table Names	189
Change Port Assignments.....	190
Configure Alert & Historical Database Connections.....	191
Troubleshoot	194
Log Files for RTView Manager.....	194
JAVA_HOME.....	194
Permissions	194
Network/DNS.....	194
Data Not Received from Data Server	195
Configure Alert Notification	195
Configure High Availability	196

Appendix A - RTView Configuration Application for Solace PubSub+ Monitor	197
Overview	198
Open the RTView Configuration Application	198
The HOME Page	199
Main Configuration Application Page	200
Saving and Applying Settings	200
General Page	201
GENERAL Tab	201
CUSTOM PROPERTIES Tab	202
Databases Page	203
Alerts Page	204
ALERTS Tab	205
HISTORY Tab	206
Security Page	207
SSL Credentials	207
Securing RTView JMX Ports	207
Secure with SSL	207
Secure with Username and Password	207
Data Server Page	209
DATA SERVER Tab	210
COLLECTOR Tab	211
Add Sender Target Dialog	212
Historian Page	213
Solution Package Configuration	214
CONNECTIONS Tab	214
Broker	215
DATA COLLECTION	216
DATA STORAGE	216
Custom Display Designer	219
Appendix B - Monitor Scripts	223
Scripts	223
rtvservers.dat	234
Appendix C - Alert Definitions	237
Alerts for Solution Package for Solace	237
Alerts for RTView Manager	245
Appendix D - Third Party Notice Requirements	247
Appendix E - Security Configuration	271
Introduction	272

Data Server	273
HTML UI	273
Data Collector	274
Configuration Application	274
Configuration Files	275
Historian.....	275
Database.....	275
Application Servers	275
RTView Manager	277
Monitored Components.....	278
Security Summary	278
Secure Installation Location - High Priority	279
Login and Servlet Authentication - High Priority	279
Application Server Security - High Priority	279
Secure Connections between RTView Processes - Medium-to-Low Priority*	279
Secure Connections to Monitored Components - Medium-to-Low Priority*	280
Secure Connections to Databases - Medium-to-Low Priority*	280
Appendix F - Limitations	281
iPad Safari Limitations	281

Preface

Welcome to the *Solace PubSub+ Monitor User's Guide*.

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

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

About This Guide

The *Solace PubSub+ Monitor 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
<i>italics</i>	Within text, new terms and emphasized words appear in italic typeface.
boldface	Within text, directory paths, file names, commands and GUI controls appear in bold typeface.
Courier	Code examples appear in Courier font: 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”](#)
- [“Documentation and Support Knowledge Base”](#)

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

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 Solution Package for Solace 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, is preconfigured as the default application server.

After you complete your evaluation, if you wish to setup and use all monitoring features in your organization, see ["Configuration"](#).

Prerequisites

- Obtain login credentials for each Solace broker you wish to monitor.
- Java JDK 1.8 64 bit
- Set the **JAVA_HOME** environment variable to point to your Java installation. For example:

UNIX

```
export JAVA_HOME=/opt/Java/jdk1.8.0
```

Windows

```
set JAVA_HOME=C:\Java\jdk1.8.0
```

- Linux Users:
 - These instructions require a Bourne-compatible shell.
 - JAVA_HOME is required to be in the PATH for Tomcat to start correctly.

To evaluate Solution Package for Solace:

1. Download **SolacePubSubMonitor_<version>.zip** to your local server and extract the files: **unzip -a SolacePubSubMonitor_<version>.zip**


Important: On UNIX systems it is a requirement that the installation directory path not contain spaces.

2. Navigate to the **SolacePubSubMonitor/bin** directory and execute **./start_servers.sh -eval** (**start_servers.bat -eval** in Windows).

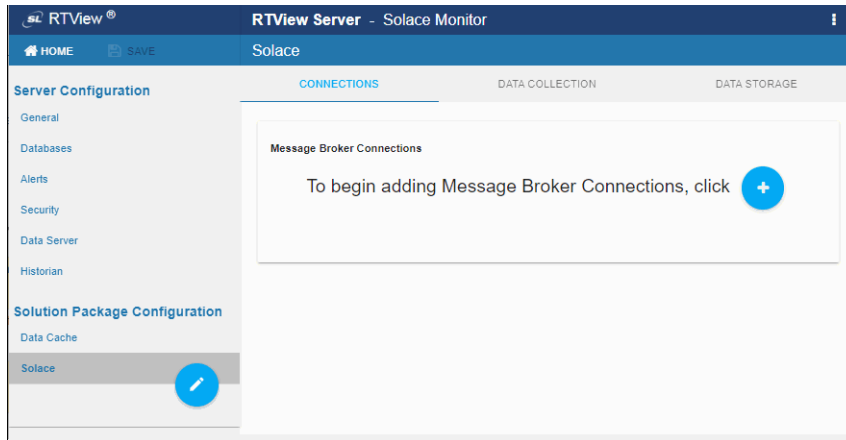
NOTE: To stop the PubSub+ Monitor when it's running in evaluation mode, execute **./start_servers.sh -eval** (**start_servers.bat -eval** in Windows).

3. Browse to the following URL and login (**rtvadmin/rtvadmin**) to open the RTView Configuration Application **HOME** page:
 - **http://IPAddress:8068/rtview-solmon-rtvadmin** if you are executing your browser on a different host than where the monitor is running.
 - **http://localhost:8068/rtview-solmon-rtvadmin** if you are executing your browser in the same host where the monitor is running.



4. Select the **Solace Monitor** project to open the **Solace** configuration page.




5. Select **Solace** in the navigation tree (left panel), then click  to add a broker connection.



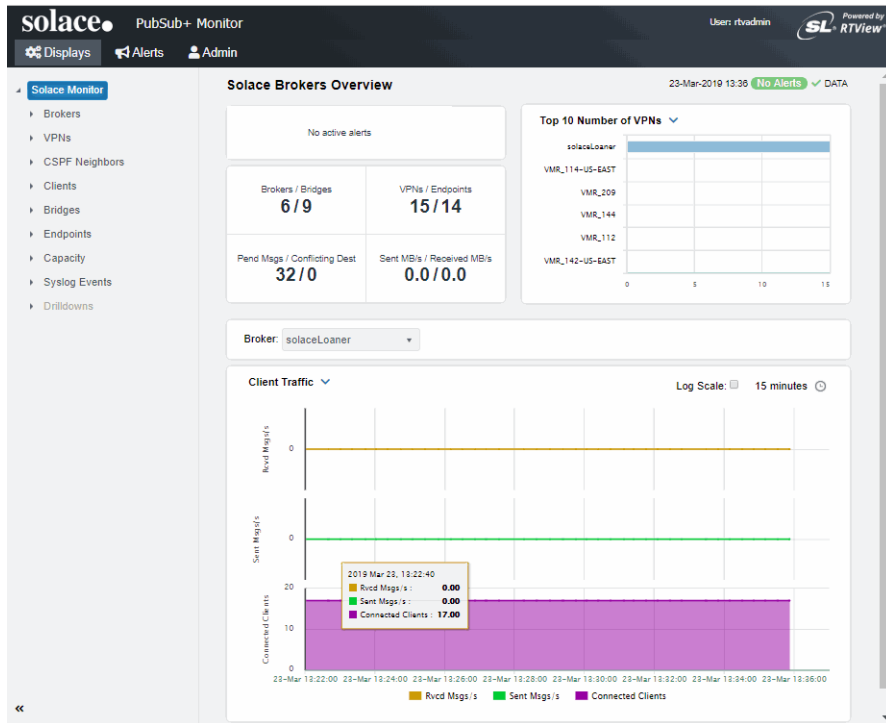
Tip: Gray text shows the default setting for the field which you can edit. To return to the default setting, delete the text you entered.

6. In the **Add Connection** dialog do the following:
 - Toggle **ON Edition** if your broker is a PubSub+ Event Broker as a service instance.
 - For Cloud Service Brokers, enter the **Connection Name**, **URL**, **Username**, **Password**, **SEMP Version*** and **VPN Name**. Turn on the **SSL Connection** toggle if your broker is SSL Secured. See ["Obtain SEMP Version"](#) for instructions about getting the SEMP version installed in your message brokers.
 - For Non-Cloud Service Brokers, enter the **Connection Name**, **Host:Port**, **Username** and **Password**. If the broker is secured, the **URL** should use **https** instead of **http**.
 - If you are connecting via SSL, toggle **ON SSL Secured**, select **SECURITY** (in the navigation tree) and fill in the **SSL Credentials** section with the appropriate **Truststore Location** and **Truststore Password** values for your Brokers.
7. Repeat these steps to add more brokers and when finished, click  to close the dialog and  (in title bar) to save your settings.

The connections you created are listed in the **Connections** tab.
8. Click  to apply changes. The data server will be available again in 10-15 seconds.
9. Browse to the Solution Package for Solace and login (username/password are **rtvadmin/rtvadmin**):

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

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



You should now see monitoring data. If you encounter issues, check the log files in the **SolacePubSubMonitor/projects/rtview-server/log** directory for errors.

You have completed the Quick Start!

10. To stop the PubSub+ Monitor, execute `./start_servers.sh -eval` (`start_servers.bat -eval` in Windows).

If you wish to setup and use all monitoring features in your organization, proceed to ["Configuration"](#).

CHAPTER 2 Introduction to the Monitor

This section contains the following:

- [“Overview,”](#) next
- [“System Requirements”](#)
- [“Install PubSub+ Monitor”](#)
- [“Upgrade PubSub+ Monitor”](#)

Overview

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

The Solution Package 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 broker versions, 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 Solution Package 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.

Solution Package for Solace is comprised of the following which you access with a browser:

- Solution Package for Solace, which monitors Solace performance metrics and used by teams to monitor the health of Solace components (brokers, bridges, clients, endpoints and VPNs).
- [“RTView Manager”](#), an application which administrators use to monitor the health of Solution Package for Solace components. That is, to monitor components of the monitoring system itself (RTView servers, JVMs, Tomcat servers, hosts and alert settings for these components). RTView Manager is installed with Solution Package for Solace and requires minimal setup.
- RTView Configuration Application, which administrators use to configure the majority of the monitoring system. For details, see [“Configuration”](#).

You can monitor Syslog events from PubSub+ message brokers using the Solace Event Module application. See [“Solace Event Module”](#) for details.

You can also install the monitor as a Solution Package (rather than a standalone product). See [“Solution Package Version”](#) for details.

Solution Package for Solace

To evaluate the Solution Package for Solace, go to [“Quick Start”](#) to get up and running with Solution Package for Solace using default settings.

Solution Package Version

The Solution Package for Solace can also be installed as a Solution Package within the RTView® Enterprise product. RTView Enterprise 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, 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, see the *RTView Enterprise 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.

Install PubSub+ Monitor

See [“Upgrade PubSub+ Monitor”](#) if you are upgrading from an earlier version of Solution Package for Solace.

To install Solution Package for Solace, download the **SolacePubSubMonitor_<version>.zip** file and unzip the **SolacePubSubMonitor_<version>.zip** file into a directory of your choosing. The **SolacePubSubMonitor/rtvapm/solmon** directory is auto-created after you unzip the file.

Important: On UNIX systems it is a requirement that the installation directory path not contain spaces.

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.

Upgrade PubSub+ Monitor

This section describes the steps necessary to upgrade existing Solution Package for Solace applications. It is organized by version. To upgrade your application, follow the steps for each version between the version you are upgrading from and the version you are upgrading to. Note that this section does not include upgrade information for the Solution Package for Solace. This section includes:

- [“Solace PubSub+ Monitor v5.2.x”](#)
- [“Solution Package for Solace v5.1.x”](#)
- [“Solution Package for Solace v5.0 and Earlier”](#) (previously referred to as “RTView Monitor for Solace”)

Solace PubSub+ Monitor v5.2.x

1. If you are upgrading from a previous release that sent SNMP notifications, you need to update the MIB in your SNMP receiver. The MIB definition in **rtvapm\common\lib\SL-RTVIEW-EM-MIB.txt** has changed to include a new field for this.
2. Some alerts are no longer used by the monitor. To remove the unused alerts from the **Alert Administration** display, remove all entries for **SolMsgRouterMsgEgressUtilHigh**, **SolMsgRouterMsgIngressUtilHigh**, **SolMsgRouterByteEgressUtilHigh** and **SolMsgRouterByteIngressUtilHigh** from the ALERTDEFS ALERTLEVELS database table. You can configure thresholds and overrides for the new alerts in the **Alert Administration** display.
3. If you enabled alert persistence, remove **SolMsgRouterMsgEgressUtilHigh**, **SolMsgRouterMsgIngressUtilHigh**, **SolMsgRouterByteEgressUtilHigh** and **SolMsgRouterByteIngressUtilHigh** from the ALERT_PERSIST_TABLE and ALERT_PERSIST_TABLE_CENTRAL to prevent errors about these alerts in the log file at startup.

Also note that if you have stored alert history, the old alert will be in the alert history. It is up to user discretion whether to keep the old alerts in history or not.

4. If you use SEMP, note that a request has been added to the supplied **pollers_sl.xml** and **groups_sl.xml** files. You may need to merge these updates with any changed versions of these files.

Note: HA connections from previous versions will automatically be split into two standalone connections. No user action is required. For details see the Release Notes.

Solution Package for Solace v5.1.x

The Solution Package for Solace file structure has been refactored. Review the release note for TN23877 for a list of all changes.

1. Download the new deliverable and extract it in a new directory on the same system as your old deliverable.
2. In the new deliverable, make a backup copy of the **projects** directory.
3. In the new deliverable, make a backup copy of the **bin** directory.
4. If you have a permanent license, copy the following file from the old installation to the new installation: **rtvapm/rtview/lib/KEYS**
5. If you modified the scripts under **RTViewSolaceMonitor\bin** in the old installation, reapply those changes to the scripts under **SolacePubSubMonitor**. Do not copy the scripts from the old installation to the new installation as they have all changed to work with the new directory structure.
6. If you modified any files under **projects\custom** in your old installation:
 - The following files under **projects\custom** were changed between 5.1 and 5.2. If you modified them in your old installation, reapply the changes to the new versions of the following files:
 - **projects/custom/src/make_classes.bat**
 - **projects/custom/src/make_classes.sh**
 - Copy any other java files you modified from **projects/custom/src/com/sl/rtvapm/custom** in the old installation to the new.
 - Execute **make_classes.bat** or **make_classes.sh** in an initialized command prompt to rebuild your custom classes against the new release.
7. The following files under **projects/rtview-server** were changed between 5.1 and 5.2. If you modified them in your old installation, reapply the changes to the new versions of these files.
 - **update_wars.bat**
 - **update_wars.sh**
8. Copy all files not mentioned in the previous step under **projects/rtview-server** from the old installation to the new installation.

9. Execute **projects/rtview-server/update_wars** in an initialized command prompt and copy the generated jars to **SolacePubSubMonitor/apache-tomcat-*-sl/webapps** or your application server.
10. The following files under **projects/rtview-manager** were changed between 5.1 and 5.2. If you modified them in your old installation, reapply the changes to the new versions of these files:
 - **update_wars.bat**
 - **update_wars.sh**
11. The following files under **projects/rtview-manager** were changed between 5.1 and 5.2. These are modified via the configuration application. Any changes you made to the **rtview-manager** via the configuration application in the previous release will be addressed in a later step.
 - **project.properties**
 - **project.properties.json**
12. Copy all files not mentioned in the previous 2 steps under **projects/rtview-manager** from old to new.
13. Execute **projects/rtview-manager/update_wars** in an initialized command prompt and copy the generated jars to **SolacePubSubMonitor/apache-tomcat-*-sl/webapps** or your application server.
14. In the previous installation, **projects/rtvservers.dat** was used for both **projects/rtview-server** and **projects/rtview-manager**. In the new installation, this has been split into **projects/rtview-server/rtvservers.dat** and **projects/rtview-manager/rtvservers.dat**. If you modified the **projects/rtvservers.dat** in your previous installation, apply the Solace monitor changes to **projects/rtview-server/rtvservers.dat** and the RTView Manager changes to **projects/rtview-manager/rtvservers.dat**.
15. If you are using the Solace Event Module:
 - Copy **rtvapm\solmon\soleventmodule\conf\soleventmodule.properties** from your old installation to your new installation to keep your old configuration. This configuration has been moved to the Solace **DATA COLLECTION** tab of the Configuration application. The next time you run the Configuration Application and save your properties files, the properties will be automatically transferred to **rtview-server/project.properties***. After that, the **soleventmodule.properties** file will no longer be used and all further configuration changes to the Solace Event Module must be done through the Configuration Application.
 - Copy the **rtvapm\solmon\soleventmodule\config\log4j2.properties** file in your old installation to **projects\rtview-server\soleventmod.log4j2.properties** in the new installation. All further changes to the Solace Event Module logging configuration should be done in **projects\rtview-server\soleventmod.log4j2.properties**.
16. If using HSQLDB, copy **projects/DATA** in your old installation to **projects/DATA** and **projects/rtview-manager/DATA** in your new installation.
17. Add new **CacheMetric** column to your alert database as described in the release note for TN24246.

18. Alter the Solace history database tables as described in the release note for TN24464.
19. If your previous installation included LDAP integration
 - Copy **apache-tomcat-*-sl/lib/ldapUser.jar** from the old installation to the new installation.
 - Copy **apache-tomcat-*-sl/conf/server.xml** from the old installation to the new installation.
 - Copy **apache-tomcat-*-sl/conf/Catalina.properties** from the old installation to the new installation.
20. If you modified **rtvapm/common/conf/sl.log4j.properties** in your old installation, copy the **sl.log4j.properties** file from your old installation to **projects/sl.log4j.properties**. All changes to the logging configuration should be made in **projects/sl.log4j.properties**.
21. Start up the new installation using the **start_servers** script under **bin**.
22. If you entered secure connection properties in the RTView Configuration Application **CUSTOM PROPERTIES** tab, they will continue to work with no changes. However, it is recommended that you remove those properties from the **CUSTOM PROPERTIES** and enter them in the Solace **CONNECTIONS** tab for easier editing in the future. The **truststore** information can be entered on the **SECURITY** tab.
23. If you made changes to the **rtview-manager** using the configuration application in your previous installation:
 - Open the configuration application at **http://localhost:3070/rtvadmin**.
 - Use **rtvadmin/rtvadmin** for the login.
 - Click **SAVE** (at the top), then **Restart Servers** to save and apply your changes.
 - Click on the **RTView Manager** server and reapply all changes you made in the previous version.
24. In the previous release, the RTView Manager was accessible at **http://localhost:3070/rtview-manager-classic** or **http://localhost:8068/rtview-manager-classic**. This has been replaced by the new HTML user interface which is available at:
http://localhost:3070/rtview-manager or **http://localhost:8068/rtview-manager**.

Solution Package for Solace v5.0 and Earlier

Users upgrading projects from the previous version must do the following:

Database Table Schemas

The database table schemas in the **rtvapm\solmon\dbconfig** directory have been updated to include all updated table schemas. To upgrade, drop the following database tables from the RTVHISTORY database:

- **SOL_BRIDGE_STATS**
- **SOL_CSPF_NEIGHBOR**
- **SOL_VPNS**
- **SOL_CLIENT_STATS**

Recreate the database tables using the appropriate table creation SQL sentence for your supported platform which are in the **rtvapm\solmon\dbconfig** directory.

RTView Configuration Application

The connection properties previously entered with the RTView Configuration Application are functional but not shown in this application correctly. To show connections properly in the RTView Configuration Application remove the past connections and recreate them in the RTView Configuration Application provided in this version of the Solace PubSub + Monitor.

CHAPTER 3 Configuration

This chapter describes how to setup Solution Package for Solace. After you complete [“Install PubSub+ Monitor”](#) instructions, you can then [“Setup Data Output Location”](#) (decide where you want to send and store your collected monitoring data).

Note: If you choose to [“Output Data to PubSub+ Monitor”](#) you can then take advantage of the many other Solution Package for Solace features in [“Optional Setup”](#).

This section contains:

- [“Setup Data Output Location”](#)
- [“Optional Setup”](#)
- [“Troubleshoot”](#)

If you wish to evaluate Solution Package for Solace using default settings, see [“Quick Start”](#).

Assumptions

This document assumes that you have:

- verified [“System Requirements”](#).
- installed the Monitor per instructions in [“Install PubSub+ Monitor”](#).

Setup Data Output Location

The first configuration step is to determine the location for the monitoring data. You have three options:

- [“Output Data to PubSub+ Monitor”](#)
- [“Output Data to InfluxDB”](#)
- [“Output Data to Solace Broker”](#)

Output Data to PubSub+ Monitor

To output monitoring data to PubSub+ Monitor, you [“Start and Login to the Solution Package for Solace”](#) and [“Configure Data Collection”](#).

Proceed to [“Start and Login to the Solution Package for Solace”](#).

Start and Login to the Solution Package for Solace

Navigate to the `SolacePubSubMonitor/bin` directory and execute the `start_servers.sh` script (or `start_servers.bat` for Windows).

Open a browser and go to:

- `http://IPAddress:8068/rtview-solmon` if you are executing your browser on a different host than where the monitor is running.

- **http://localhost:8068/rtview-solmon** if you are executing your browser in the same host where the monitor is running.

Use **rtvadmin/rtvadmin** for username/password.

The Solution Package for Solace opens. The displays populate with data after you add connection properties for your Solace Message Brokers (which is subsequently described in these instructions).

Proceed to [“Open the RTView Configuration Application”](#).


Open the RTView Configuration Application

Open a browser and go to:

- **http://IPAddress:8068/rtview-solmon-rtvadmin** if you are executing your browser on a different host than where the monitor is running.
- **http://localhost:8068/rtview-solmon-rtvadmin** if you are executing your browser in the same host where the monitor is running.

Use **rtvadmin/rtvadmin** for username/password.


The RTView Configuration Application **HOME** page opens.



Select the **Solace Monitor** project.

The main configuration page for the **RTView Server - Solace Monitor** project opens.

The navigation tree is in the left panel and the **General** and **Custom Properties** tabs are shown in the upper part of the main page. The name of the selected tab is highlighted and the other tabs are grayed out. You click on either of the grayed tabs to change the selected tab.



These instructions use the following format to describe navigation to each tab: **Navigation tree>Tab**. For example, the figure above illustrates the **General>GENERAL Tab**.

Proceed to [“Configure Data Collection”](#) (a required configuration).

Configure Data Collection

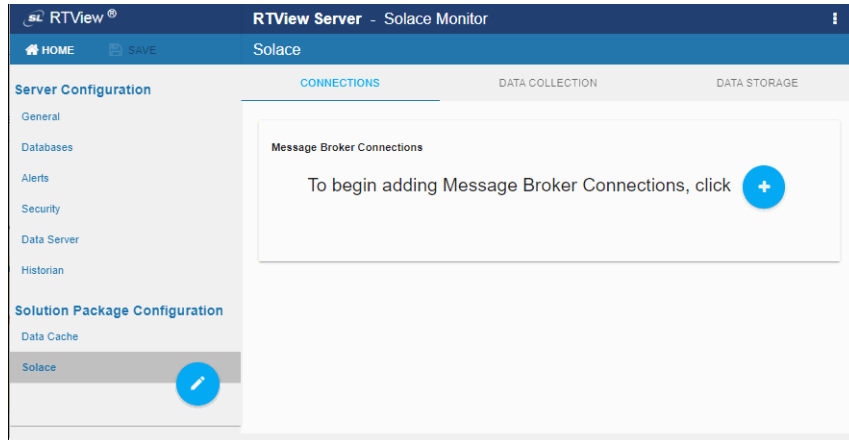
This section describes how to define the connection for the brokers you wish to monitor and verify that data is collected from them. This configuration must be performed before running any deployment of the Monitor. This configuration is the only required configuration.

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.

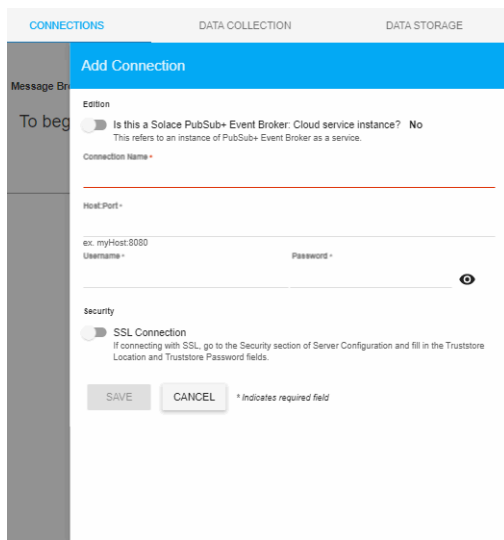
Note that for Solace Cloud Event Brokers you will need the exact SEMP version on each of your Solace Cloud Brokers. See [“Obtain SEMP Version”](#).

To define Solace Broker connections:

1. "Open the RTView Configuration Application", select **Solace** (in navigation tree) > **CONNECTIONS** tab and click .



The **Add Connection** dialog opens.



2. In the **Add Connection** dialog do the following:

- Toggle **ON Edition** if your broker is a PubSub+ Event Broker:Cloud service instance.
- For Cloud Service Brokers, enter the **Connection Name**, **URL**, **Username**, **Password**, **SEMP Version*** and **VPN Name**. Turn on the **SSL Connection** toggle if your broker is SSL Secured. See "[Obtain SEMP Version](#)" for instructions about getting the SEMP version installed in your message brokers.
- For Non-Cloud Service Brokers, enter the **Connection Name**, **Host:Port**, **Username** and **Password**.
- If you are connecting via SSL, toggle **ON SSL Secured**, select **SECURITY** (in the navigation tree) and fill in the **SSL Credentials** section with the appropriate **Truststore Location** and **Truststore Password** values for your Brokers.

- Repeat these steps to add more brokers and when finished, click **SAVE** to close the dialog and **SAVE** (in title bar) to save your settings.

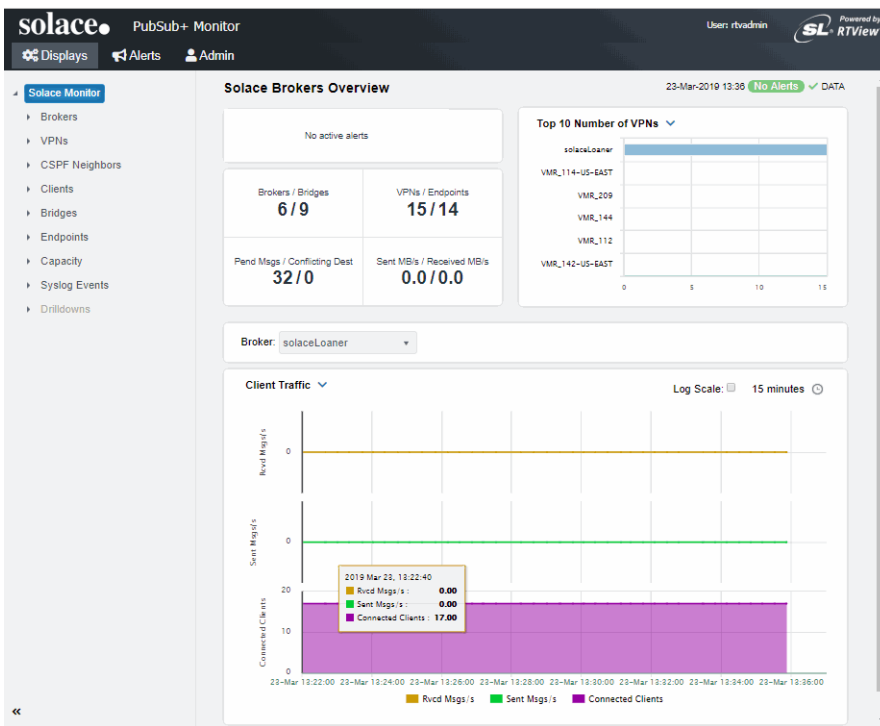
The connections you created are listed in the **Connections** tab. For HA pairs, the connection string for the backup broker has "-B" concatenated to it.

- Click **RESTART SERVERS** to apply changes. The data server will be available again in 10-15 seconds.
- Open a browser and go to the Solution Package for Solace:

http://IPAddress:8068/rtview-solmon if you are executing your browser on a different host than where the monitor is running.

http://localhost:8068/rtview-solmon if you are executing your browser in the same host where the monitor is running.

Use **rtvadmin/rtvadmin** for username/password.



You should now see monitoring data. If you encounter issues, check the log files in the **SolacePubSubMonitor/projects/rtview-server/log** directory for errors.

- If you are going to:
 - "Output Data to InfluxDB", proceed to "Create Database and User Account".
 - "Output Data to PubSub+ Monitor", congrats! You have finished the required setup. You can take advantage of the many optional Solution Package for Solace features. See "Optional Setup" for details.

Obtain SEMP Version

This section only applies to Solace Cloud Event Brokers. You need to provide the exact SEMP version on each of your Solace Cloud Brokers.

Use the Solace Cloud console to get the value for the **Solace Broker Version** field, under **Stats**. The broker version aligns with the SEMP v1 version. You will use the first three digits, including any decimal points, of the value shown in the Solace Cloud Event Broker **Solace Broker Version** field, concatenated with this string: **VMR**.

For example, if the value for the **Solace Broker Version** field is **9.1.1.1.0**, you enter: **9.1.1VMR**

Output Data to InfluxDB

To output monitoring data to InfluxDB, you need to [“Start and Login to the Solution Package for Solace”](#), [“Open the RTView Configuration Application”](#) to configure connections to Solace Brokers and then [“Create Database and User Account”](#).

Create Database and User Account

This section assumes you have already configured connections to your Solace Brokers using the RTView Configuration Application.

To send monitoring data to InfluxDB you edit the **stats-receiver.properties** file, located in the **Data Collector** directory, define the metrics to poll and store, and specify the IP address and port number on the InfluxDB platform.

1. Create a database and a user with read and write privileges in your InfluxDB platform using the following InfluxDB shell commands:

```
create database <yourSolaceStatsDB>
```

```
use <yourSolaceStatsDB>
```

```
create user <SolaceStatsUser> with password <'yourPwdWithSingleQuotes'>
```

```
with all privileges grant all on <yourSolaceStatsDB> to <SolaceStatsUser> exit
```

2. Edit the **stats-receiver.properties** file, located in the directory where the data server will be started.

3. Enable InfluxDB Tap by commenting out the following line:

```
#TAP_PLUGIN_CLASS = com.sl.statsds.RTViewStatsTap
```

and uncommenting the following line:

```
TAP_PLUGIN_CLASS =
```

```
com.solace.psg.enterprisestats.receiver.influxdb.InfluxDBStatsTap
```

This enables the plugin to send monitoring data to InfluxDB.

Note: At present, there is no option to send monitoring data to both InfluxDB and PubSub+ Monitor. To revert back to collect all caches and send data to the PubSub+ Monitor, uncomment the **RTViewStatsTap** line and comment out the **InfluxDBStatsTap** line. Also revert **DB_FIELD_SUBSCRIPTIONS** to its initial value (**DB_FIELD_SUBSCRIPTIONS =** which means the filter passes everything, as appropriate for the PubSub+ Monitor).

4. Select the list of topics identifying the fields to write to InfluxDB by copying the list of topics in the **DB_FIELD_SUBSCRIPTIONS** property. These topics are semicolon separated and

multiple lines are identified by "\" at the end of the topic. Initially you can choose from the examples that are provided in the **stats-receiver.properties** file itself. For example:

```
DB_FIELD_SUBSCRIPTIONS=\
SYSTEM_CONFIG-SYNC/authentication/client-certificate/max-certificate-chain-depth;\
SYSTEM_MEMORY/subscription-memory-usage-percent;\
SYSTEM_MEMORY/physical-memory-usage-percent;\
SYSTEM_MEMORY/slot-infos/*/nab-buffer-load-factor;\
SYSTEM_STATS_CLIENT/>;\
SYSTEM_STATS_NEIGHBOR/>;\
SYSTEM_CSPF_NEIGHBOR_STATS/>;\
SYSTEM_MSG-SPOOL_DETAIL/*;\
SYSTEM_MSG-SPOOL_STATS/>;\
VPN_STATS/maximum-spool-usage-mb;\
>;\
```

Refer to Solace documentation for the complete list of metrics that can be requested.

5. Add the connection properties to InfluxDB by scrolling down to the **Influx DB Properties** and defining the following:

- **INFLUXDB_HOST**: the hostname and port used for Influx DB. For example, if you run it locally in the default port: **INFLUXDB_HOST = localhost:8086**.
- **INFLUXDB_DB**: the database name to which InfluxDB will write. For example: **INFLUX_DB=yourSolaceStatsDB**
- **INFLUXDB_USER**: the user previously created in Influx DB to execute the inserts in the DB. For example: **INFLUXDB_USER = SolaceStatsUser**
- **INFLUXDB_PASSWORD**: the password you set in InfluxDB encrypted with the Solace PSG Password Utility. To use the Solace PSG Password Utility:
 - Open a command prompt or terminal and initialize it
 - Change directory (**cd**) to **<installation_dir>/rtvapm>** and execute **./rtvapm_init.sh (UNIX)**

or

cd <installation_dir>\rtvapm> and execute **rtvapm_init.bat (Windows)**

- Change directory (**cd**) to **rtvapm\solmon\bin** and execute:

pwd-utility[.bat] <yourPwdString>

The following message appears:

"The encrypted password is: 'encryptedString' (without the quotes).

You should use this value in the password field in configuration property files."

You should now see monitoring data being stored in the InfluxDB database. If you encounter issues, check the log files in the **SolacePubSubMonitor/projects/rtview-server/log** directory for errors and verify that InfluxDB is available and running.

You have completed configuring data collection! There are no other required configurations.

Output Data to Solace Broker

To send monitoring data to a Solace Broker do the following. The files you edit are located in the **rtvapm/solmon/bin/config** directory. No other configuration steps than those provided here are needed to output monitoring data to a Solace Broker. Refer to Solace documentation for additional information about the available message formats you can choose for sending the data.

1. Define the Solace Brokers to monitor by editing the **<primary> ... </primary>** section of the **appliance_config_demo.xml** file (this adds connection properties to the monitored Solace Brokers). Add as many **<primary> ... </primary>** sections as brokers you want to monitor.
2. Define the Solace Broker to receive monitoring data by editing the **<mgmt.-msg-bus> ... </mgmt.-msg-bus>** section from the **appliance_config_demo.xml** file. You should define the message format for the data being transmitted by choosing one of the container factories: **SempXmlFragmentFactory**, **JsonMapFactory** or **StdMapFactory** for SEMP, and **JSON** or standard message formats respectively.
3. If the data to be polled isn't already defined in the preconfigured XML files, or you need additional poller groups for different monitoring options, do the following:
 - Define the monitoring data to poll by editing the **pollers_sl.xml** file (this file contains SEMP request details and response parsing specifics which the poller sends to Solace Brokers).
 - Define the poller groups that you want to use by editing the **groups_sl.xml** file (this file enables you to separate published statistics into groups of interest, publishes the statistics on the associated topic, and provides the configured poll interval).
4. Start StatsPump as follows:

In a Windows command prompt or UNIX terminal, go to the **SolacePubSubMonitor/rtvapm** directory and execute **rtvapm_init.bat** (Windows) or **rtvapm_init.sh** (UNIX).

Change directory (**cd**) to **rtvapm/solmon/bin** directory and execute the following in the order provided (if you change the order it will not execute properly):

```
[statspump|statspump.bat] config\pollers_sl.xml config\groups_sl.xml
config\appliance_config_demo.xml
```

You should now see published monitoring data in the receiving Solace Broker.

You have finished configuration instructions to send monitoring data to a Solace Broker!

Optional Setup

This section describes how to setup optional features that are available if you [“Output Data to PubSub+ Monitor”](#):

- [“Choose and Setup an Application Server”](#): You have two options:
 - [“Using the Pre-configured Apache Tomcat”](#): This option allows you to optionally [“Integrate LDAP with Tomcat and HTML UI”](#).

- [“Using an Alternate Application Server”](#)
- [“Modify Default Polling Rates for Solace Caches”](#): Change the default polling rates for Solace caches.
- [“Modify Default Settings for Storing Historical Data”](#): Change the default settings for how historical data is collected, aggregated and stored in caches.
- [“Change Port Assignments”](#): Change the default port settings.
- [“Configure Alert & Historical Database Connections”](#): Configure a production database.
- [“Configure Alert Notification”](#): Configure alerts to execute an automated action (for example, to send an email alert).
- [“High Availability”](#): Configure failover and failback for Data Servers and the Historian.
- [“Configure PubSub+ Message Broker & Syslog Destination”](#): Configure PubSub+ Brokers and Solace Monitor to receive Syslog events and activate Syslog event-driven alerts.
- [“Troubleshoot”](#): Investigate configuration issues.

Choose and Setup an Application Server

Solution Package for Solace requires an application server. You have two options:

- [“Using the Pre-configured Apache Tomcat”](#): This option allows you to [“Integrate LDAP with Tomcat and HTML UI”](#).
- [“Using an Alternate Application Server”](#)

Using the Pre-configured Apache Tomcat

Solution Package for Solace includes a pre-configured Apache Tomcat installation which hosts all of the servlets necessary to run the Monitor on port **8068**. If you would like to use this application server for your deployment, no further configuration is required. You can optionally change user names and passwords for the servlets hosted in Tomcat in

SolacePubSubMonitor/apache-tomcat*/sl/conf/tomcat-users.xml. The user names and passwords in this file can be changed, but you must assign them one of the defined roles as these are required by the servlets. For details about predefined user roles, see [“User Permissions”](#).

You can optionally [“Integrate LDAP with Tomcat and HTML UI,”](#) next.

Proceed to [“Start and Login to the Solution Package for Solace”](#).

Integrate LDAP with Tomcat and HTML UI

This section describes how to setup Tomcat so that you can log into RTView applications using your LDAP username and password and automatically associate that with an RTView role (for example, **rtvadmin**, **rtvuser**, and so forth). This feature uses Apache Tomcat’s built-in LDAP integration, the JNDI Realm.

To integrate LDAP with Tomcat and the HTML UI you edit two configuration files. Then, when a user logs into RTView, it is located in the directory and authenticated using the given password. Users can be located by a direct name-to-DN association or by searching for a specified attribute. You can search using the credentials of the given user or supply a **connectionName** and **connectionPassword** and search using those credentials.

After a user is authenticated, they are associated with an RTView role according to the entries in the **rtview-roles.txt** file. Users given the role:

- **rtvuser** will be able to log into the HTML UI.
- **rtvadmin** will be able to log into the Configuration Application.
- **rtvalertmgr** will be able to manage alerts.

Note that you must restart the Tomcat server after editing **server.xml** but it is not necessary to restart it after editing **rtview-roles.txt**.

To setup LDAP integration:

1. Open the **server.xml** file, located in the **TOMCAT_HOME/conf** directory.
2. Comment out this section:

**<Realm className="org.apache.catalina.realm.UserDatabaseRealm" and
uncomment the section <Realm
className="org.apache.catalina.realm.RTViewJNDIRealm"**

The result should look like this:

```

1      <!--
2      <Realm className="org.apache.catalina.realm.UserDatabaseRealm"
3          resourceName="UserDatabase"/>
4      -->
5      <Realm className="org.apache.catalina.realm.RTViewJNDIRealm"
6          connectionURL="ldap://localhost:389"
7          />

```

3. Specify your LDAP server connection information and other LDAP options.
4. Restart the Tomcat server.
5. Open the **rtview-roles.txt** file, located in the **TOMCAT_HOME/conf** directory, to specify your role-to-attribute mappings.

Refer to Tomcat JNDI Realm documentation for additional information and the following examples below.

Example 1

You have four users in a container named **Users**. You want to find them by name and associate each one with RTView roles. Tom needs the **rtvadmin** role. Nancy needs both the **rtvadmin** and **rtvalertmgr** roles. Joe and Susan need the **rtvuser** role. To set these up you edit the files as follows:

- In **server.xml** you include:

userPattern="CN={0},CN=Users,DC=mycompany,DC=com"

- In **rtview-roles.txt** you include:

rtvadmin:CN=Tom,CN=Users,DC=mycompany,DC=com

rtvadmin,rtvalertmgr:CN=Nancy,CN=Users,DC=mycompany,DC=com

rtvuser:CN=Joe,CN=Users,DC=mycompany,DC=com

rtvuser:CN=Susan,CN=Users,DC=mycompany,DC=com

Example 2

Your groups are in a container named **Roles**. You want any members of the **Administrators** group to get the **rtvadmin** and **rtvuser** roles. You also want any members of the **Users** group to get the **rtvuser** role. To set these up you edit the files as follows:

- In **server.xml** you include:

```
userPattern="CN={ 0 },CN=Users,DC=mycompany,DC=com"
```

```
userRoleName="memberOf"
```

- In **rtview-roles.txt** you include:

```
rtvadmin,rtvuser:CN=Administrators,CN=Roles,DC=mycompany,DC=com
```

```
rtvuser:CN=Users,CN=Roles,DC=mycompany,DC=com
```

For additional details, refer to Apache Tomcat JNDI Realm documentation.

Proceed to [“Start and Login to the Solution Package for Solace”](#).

Using an Alternate Application Server

Alternately, you can use another application server. To use another application server:

1. In a windows command prompt or UNIX terminal go to **SolacePubSubMonitor/rtvapm** and execute **rtvapm_init.bat** (Windows) or **.rtvapm_init.bat** (UNIX).
2. Change directory (**cd**) to **SolacePubSubMonitor/projects/rtview-server** and execute **update_wars.bat** (Windows) or **update_wars.sh** (UNIX).
3. Deploy the resulting war files to your application server.
4. If you are using the RTView Manager, repeat the previous two steps in the **projects/rtview-manager** directory.
5. Add the following user roles to your application server: **rtvuser**, **rtvadmin**, **rtvalertmgr**. For details about user roles, see [“User Permissions”](#).

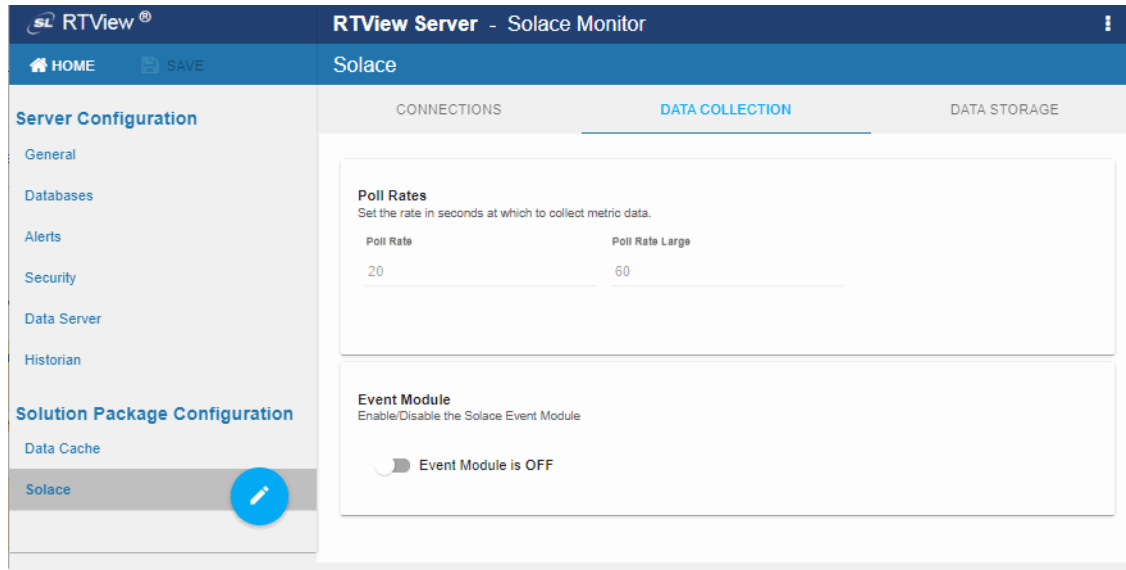
The instructions in this document refer to the pre-configured Apache Tomcat host and port (**localhost:8068**). When following instructions, use the application server's host and port instead.

Proceed to [“Start and Login to the Solution Package for Solace”](#).

Modify Default Polling Rates for Solace Caches

To modify the default polling rate settings for Solace caches, perform the following:

“Open the [RTView Configuration Application](#)” and go to **Solace>DATA COLLECTION** tab.



Poll Rate: Collection period in seconds. This configuration element affects the following caches: SolEndpointStats, SolEndpoints, SolClients, SolClientStats, SolBridges, SolAppliances, SolBridgeStats, SolApplianceInterfaces and SolApplianceMessageSpool.

Poll Rate Large: Slower collection period in seconds for monitoring data that can impact the performance of the monitoring systems if the rate is very fast. This configuration element affects the following caches: SolCspfNeighbors, SolAppliances and SolEnvironmentSensors.

Solace Event Module Alerts Clear Time: Defines the time interval, in seconds, when non-clearable event alerts from the Solace Event Module will be dismissed from the monitor.

- **SAVE** your settings, then click **RESTART SERVERS** to apply changes. The data server will be available again in 10-15 seconds.

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 to modify the default behavior of the data being collected, aggregated and stored.

- “[Define the Storage of In Memory History](#)”: Specify the maximum number of history rows to store in memory.
- “[Define Compaction Rules](#)”: Define rules for reducing the amount of data stored over time.
- “[Define Duration](#)”: Specify when data becomes expired and/or deleted from the Monitor.
- “[Enable/Disable Storage of Historical Data](#)”: Choose the metrics you want to store in the database and specify a prefix for history table names.

- [“Define Prefix for All History Table Names”](#): Specify a prefix to prepend to database table names.



Define the Storage of In Memory History

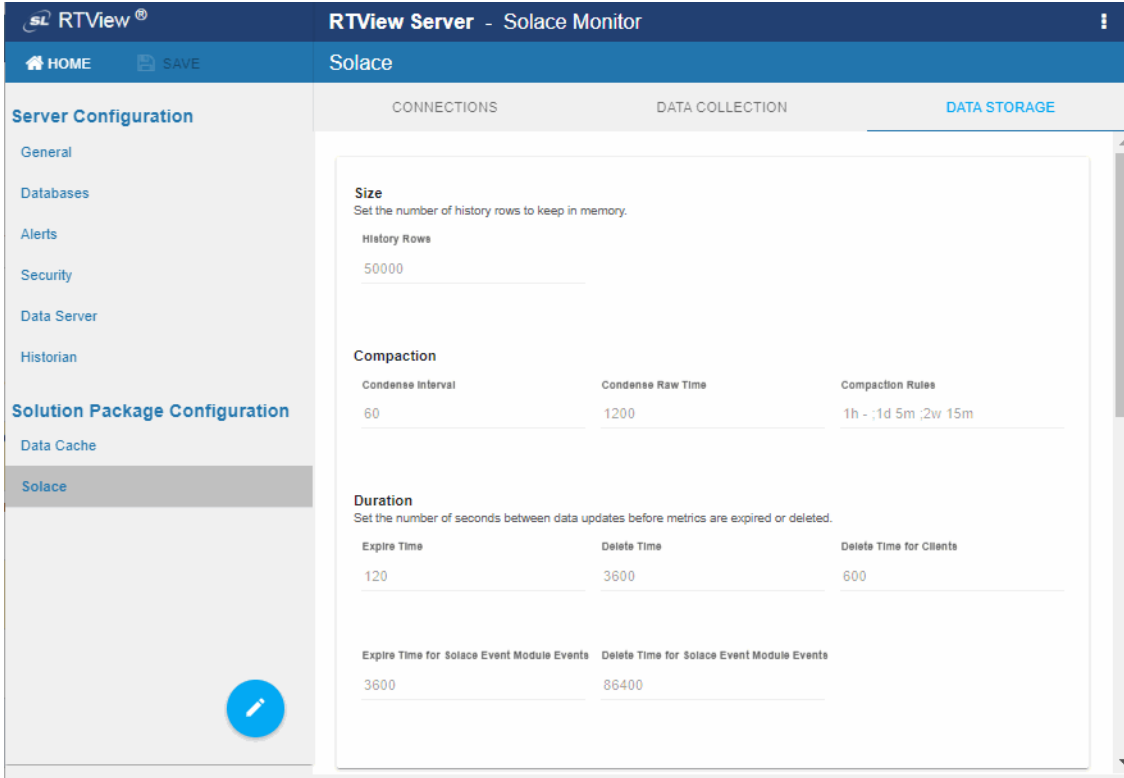
You can define the maximum number of history rows to store in memory in the **Solace/Data Storage/History Rows** property. This property can improve Monitor responsiveness.

Note that changing this value is only recommended if you have a high degree of understanding about how historical data is being stored in memory, as well as how that data is compacted and stored in the database.

The **History Rows** property defines the maximum number of rows to store for the SolVpns, SolClientStats, SolAppliances, SolEndpoints, SolCspfNeighbors, SolBridgeStats, SolApplianceInterfaces, SolApplianceMessageSpool, SolEndpointStats, SolEventModuleEvents and SolAppliancesQuality caches. The default setting for **History Rows** is **50,000**.

To modify the default settings:

- [“Open the RTView Configuration Application”](#) and go to **Solace>DATA STORAGE** tab.
- Under **Size**, enter the desired number of rows in the **History Rows** field.
-  **SAVE** your settings, then click  **RESTART SERVERS** to apply changes. The data server will be available again in 10-15 seconds.



RTView Server - Solace Monitor

Solace

DATA STORAGE

Size
Set the number of history rows to keep in memory.

History Rows
50000

Compaction

Condense Interval	Condense Raw Time	Compaction Rules
60	1200	1h - ;1d 5m ;2w 15m

Duration
Set the number of seconds between data updates before metrics are expired or deleted.

Expire Time	Delete Time	Delete Time for Clients
120	3600	600



Expire Time for Solace Event Module Events: 3600
Delete Time for Solace Event Module Events: 86400

Define Compaction Rules

Data compaction, essentially, is reducing redundancy in the data to be stored in the database by using a rule so that you store sampled data instead of raw data, which prevents storing of redundant data which potentially can overload the database. The compaction rule is defined through the following fields:

- **Condense Interval:** The time interval at which the cache history is condensed. The default is **60** seconds, which means that every **60** seconds all rows of the same index are condensed. As a result of this first condensing operation there will be only one row per index every minute. 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 memory. 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 are aggregated by averaging rows with the following rule **1h -;1d 5m;2w 15m**, which means the data from the last hour is not aggregated (1h - rule), the data from the last day is aggregated every 5 minutes (1d 5m rule), and the data from the last 2 weeks old is 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 Configuration Application”](#) and go to **Solace>DATA STORAGE** tab.
- Under **Compaction**, enter values in the **Condense Interval**, **Condense Raw Time** and **Compaction Rules** fields.
-  your settings, then click  to apply changes.



Define Duration

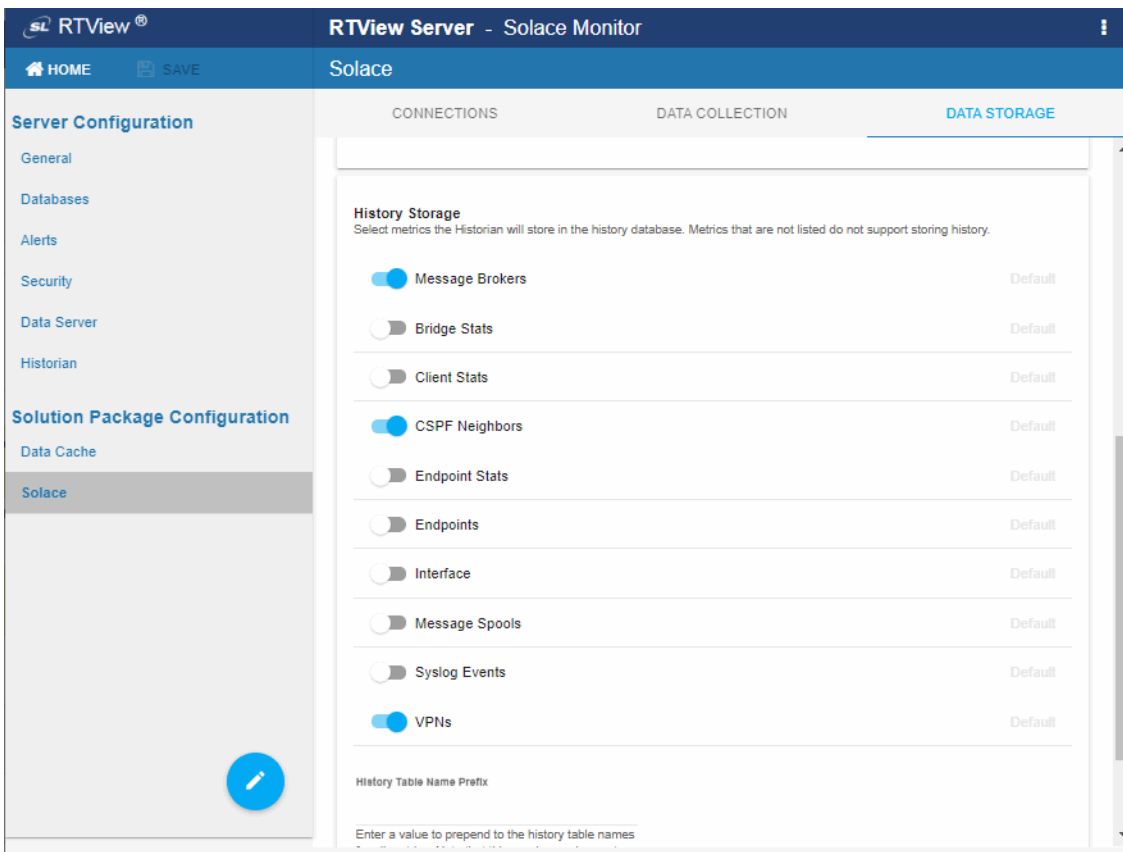
The data for each metric is stored in a specific cache and, when the data is not updated in a certain period of time, that data either marked as expired or, if it has been expired over an extended period of time, it is deleted from the cache altogether.

- **Expire Time:** This field sets the period of time when the Expire metric from the cache is set to true indicating the entry row is expired. The default expiration time is 120 seconds. The following caches have this attribute defined: SolVpns, SolBridges, SolClients, SolClientStats, SolAppliances, SolEndpoints, SolCspfNeighbors, SolBridgeStats, SolApplianceInterfaces, SolApplianceMessageSpool, SolEndpointStats, SolEnvironmentSensors and SolAppliancesQuality.
- **Delete Time:** This field sets the period of time that a given entry row should be expired before it gets deleted from the cache. It defaults to 3600 seconds and applies to the following caches: SolVpns, SolBridges, SolEndpoints, SolBridgeStats, SolEndpointStat and SolEnvironmentSensors caches.
- **Delete Time for Clients:** The meaning of this field is the same as of the Delete Time but applying to the SolClients and SolClientStats caches. The default is 600 seconds.
- **Expire Time for Solace Event Module Events:** This field sets the expiration period exclusively for the SolEventModuleEvents cache, which defaults to 3600 seconds.
- **Delete Time for Solace Event Module Events:** The meaning is as the two previous fields but for the SolEventModuleEvents cache. The default is 1 day (86,400 seconds).

Enable/Disable Storage of Historical Data

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:

- “Open the RTView Configuration Application” and go to **Solace>DATA STORAGE** tab.
- 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 Brokers), SolBridgeStats (Bridge Stats), SolClientStats (Client Stats), SolCspfNeighbors (CSPF Neighbors), SolEndpointStats (Endpoint Stats), SolEndpoints (Endpoints), SolApplianceInterfaces (Interface), SolApplianceMessageSpool (Message Spools), SolEventModuleEvents (Syslog Events) and SolVpns (VPNs).
-  **SAVE** your settings, then click  **RESTART SERVERS** to apply changes.





Define Prefix for All History Table Names

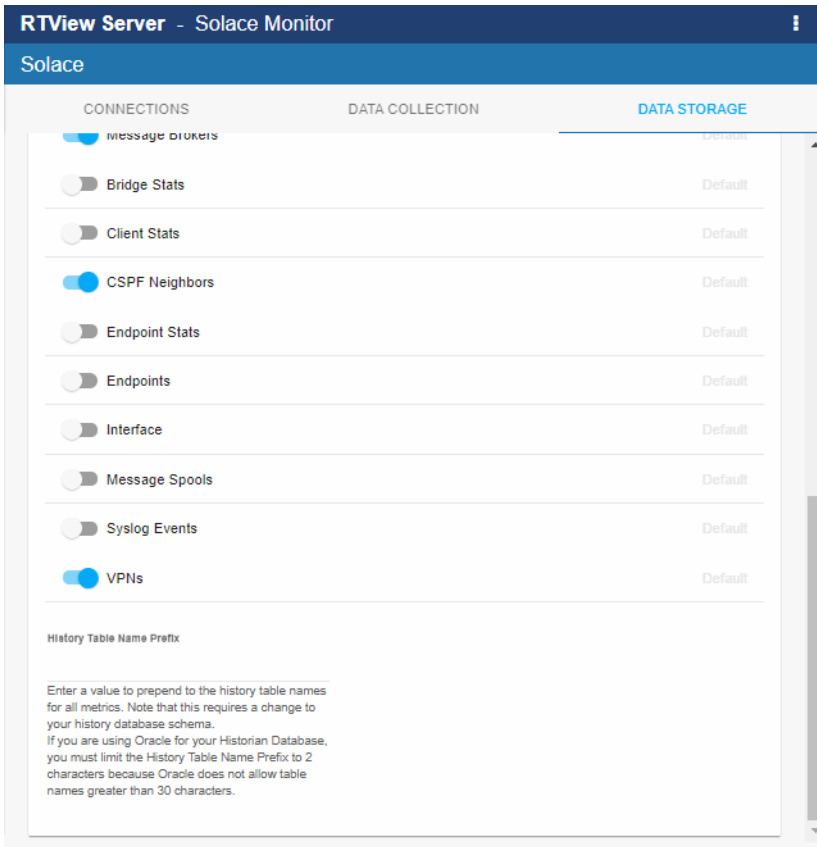
The **History Table Name Prefix** field allows you to define a prefix that is added to the database table names so that the Monitor can differentiate history data between data servers when you have multiple data servers with corresponding Historians using the same solution package(s) and database. In 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 need to create the corresponding tables in your database as follows:

- Locate the .sql template for your database under **SolacePubSubMonitor/rtvapi/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 Configuration Application](#)”, go to **Solace>DATA STORAGE** tab and scroll down to the bottom of the page.
- In the **History Table Name Prefix** field, enter the desired prefix name.
-  **SAVE** your settings, then click  **RESTART SERVERS** to apply changes.





Change Port Assignments

This configuration is optional.

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
Collector 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

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

1. ["Open the RTView Configuration Application"](#) and go to **General>GENERAL** tab.
2. Under **Ports** (scroll down to the bottom of the page), 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  (in the title bar), then click  to apply changes.
4. Edit the **update_wars** (.bat or .sh) file and change the port prefix for all ports to the prefix you just specified.
5. Rebuild the war files and install them to the application server by executing the following script, located in the **SolacePubSubMonitor/bin** directory:
Windows:
make_all.bat
UNIX:
./make_all.sh

Configure Alert & Historical Database Connections

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*.

This section describes how to setup an alternate production database, and how to configure the Alert Settings Database connection and the Historian Database connection. You connect and configure the databases using the RTView Configuration Application. You also copy portions of the **database.properties** template file (located in the **common\dbconfig** directory) into the RTView Configuration Application.

Monitor Databases

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 cannot work correctly.

Historian

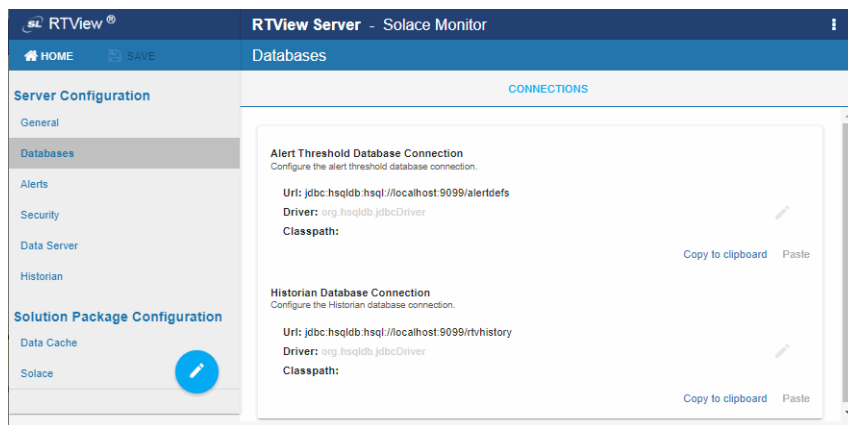
The RTVHISTORY database contains the historical monitoring data to track system behavior for future analysis, and to show historical data in displays.

To Configure the ALERTDEFS and RTVHISTORY Databases:

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

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

2. Open the **database.properties** template file, which is located in the **common\dbconfig** directory, find the line that corresponds to your supported database in the "Define the ALERTDEFS DB" section and make a note of this information. Keep the **database.properties** template file open.
3. "Open the RTView Configuration Application" and go to **Databases>CONNECTIONS** tab.



4. Click the **Alert Threshold Database Connection** to open the **Edit Connection** dialog.
5. Enter the information (you previously noted from the **database.properties** file) 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.

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

6. Return to the **database.properties** template file, which is located in the **common\dbconfig** directory, find the line that corresponds to your supported database in the "Define the RTVHISTORY DB" section and make a note of this information.

7. In the RTView Configuration Application, click the **Historian Database Connection** to open the **Edit Connection** dialog.

8. Enter the information (you previously retrieved from the **database.properties** file) 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.

9. Click  to store the newly added connection and close the dialog and  (in title bar) to save your settings.

10. Click  to apply changes.

11. 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, where:

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

- **Alert Settings**

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

- **Historian**

```
SolacePubSubMonitor/rtvapm/solmon/dbconfig/  
create_solmon_history_tables_<db>.sql  
  
SolacePubSubMonitor/rtvapm/common/dbconfig/  
create_common_history_tables_<db>.sql  
  
SolacePubSubMonitor/rtvapm/rtvmgr/dbconfig/  
create_rtvmgr_history_tables_<db>.sql
```

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.

Troubleshoot

This section includes:

- ["Log Files for Solace"](#)
- ["JAVA_HOME"](#)
- ["Permissions"](#)
- ["Network/DNS"](#)
- ["Data Not Received from Data Server"](#)
- ["Obtain SEMP Schemas"](#)

Log Files for Solace

When any Solution Package 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.

Solution Package for Solace Log Files

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

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

RTView Manager Log Files

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

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

JAVA_HOME

If you encounter issues starting Solution Package for Solace 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 confirm with your network administrator that you're not being blocked from accessing the remote system.

Data Not Received from Data Server

In the Solution Package 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. Do the following:

1. ["Open the RTView Configuration Application"](#) and go to the **Solace>CONNECTIONS** tab.
2. Verify the connection parameters associated with your brokers.
3. Verify the SEMP version is correct for each of your Cloud Brokers (monitoring data cannot be collected if the SEMP version is incorrect) and make corrections if necessary.

Click  **SAVE** in the title bar when finished, then click  **RESTART SERVERS** to apply changes. It takes about 10-15 seconds for the data server to be available again.

4. In the Solution Package for Solace, go to the **Admin>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).

Obtain SEMP Schemas

When SEMP schemas that are used for connecting to a Solace Broker are missing, the Broker is not shown in the PubSub+ Monitor **Broker Table** and the log from the dataserver under the **projects\rtview-server\logs** directory shows the following exception:

```
... java.lang.IllegalStateException: Have not loaded schemas for 'soltr/9_OVMR'.  
Ensure schema version looks like 'soltr/x_y'. Call loadSchemas() first, or import  
new schema files...
```

To resolve this problem, download the schemas from the Solace Customer Portal (<https://products.solace.com/>) where you as a customer have access to download Solace products as well as the SEMP schemas from any supported release (either Software or Appliance Brokers).

Be aware that the SEMP schemas from Software and Appliance Brokers, even from the same version, might differ. Therefore, verify that the downloaded files for either Software or Appliance Brokers are uniquely identified.

After you download the SEMP schemas do the following to include them:

1. Change directory (**cd**) to the **resources** directory (for example, **yourProjectDir/rtvapm/solmon/lib/ext/resources**) and create a separate directory for the downloaded SEMP schemas.
2. Copy the two schema files into the newly created directory.
3. **Stop/start** PubSub+ Monitor and verify that the missing Broker is connected and data is being collected properly.
4. Verify that the **dataserver.log** file no longer shows the missing schema error.

Contact SL Technical Support if you have issues downloading or adding these files to the product.

CHAPTER 4 Additional Configurations

This section contains the following:

- [“Solace Event Module”](#)
- [“High Availability”](#)
- [“Property Editor REST API”](#)
- [“Create Custom Alerts”](#)

Solace Event Module

You can monitor Solace PubSub+ message broker Syslog events using the Solace Event Module application. To use the Solace Event Module you [“Configure PubSub+ Message Broker & Syslog Destination”](#) to send Syslog messages, and set the Solace Event Module to run and listen for the Syslog messages the broker sends.

This section contains:

- [“Introduction”](#)
- [“Configure PubSub+ Message Broker & Syslog Destination”](#)
- [“Configure Solace Event Module”](#)
- [“Solace Event Module Logging”](#)

Introduction

To monitor the PubSub+ message brokers using Syslog events, you can use the Solace Event Module application. The Solace Event Module listens for Syslog event messages that are generated by Solace PubSub+ message brokers and filters them to generate Syslog event-based alerts when required.

The Syslog event messages generated in the PubSub+ message brokers are forwarded to the SolEventModuleEvents cache from the RTView Solace Data Server. The events that trigger alerts are stored in the SolEventModuleAlerts cache from the RTView Solace Data Server.

The Solace Event Module is licensed separately from the Solution Package for Solace. And therefore, it is not executed by default and requires additional configuration.

Configure PubSub+ Message Broker & Syslog Destination

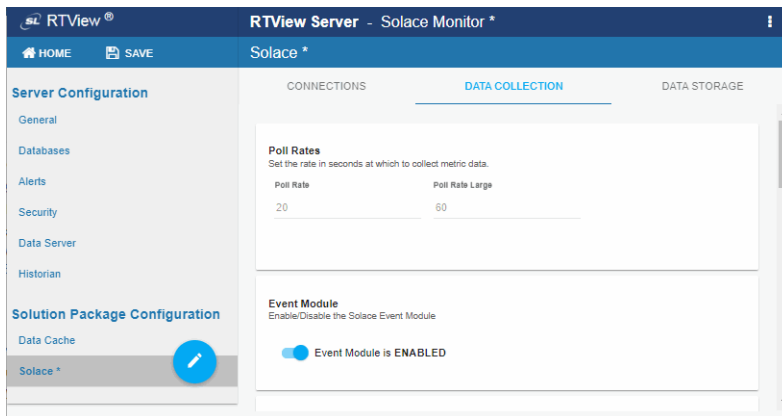
To use the Solace Event Module, you must first configure your Solace brokers to send Syslog messages for either the **system.log** or **event.log** and also configure a receiver for those messages that can be accessed by your Solace data server.

For the configuration of PubSub+ message brokers with Syslog, please refer to:

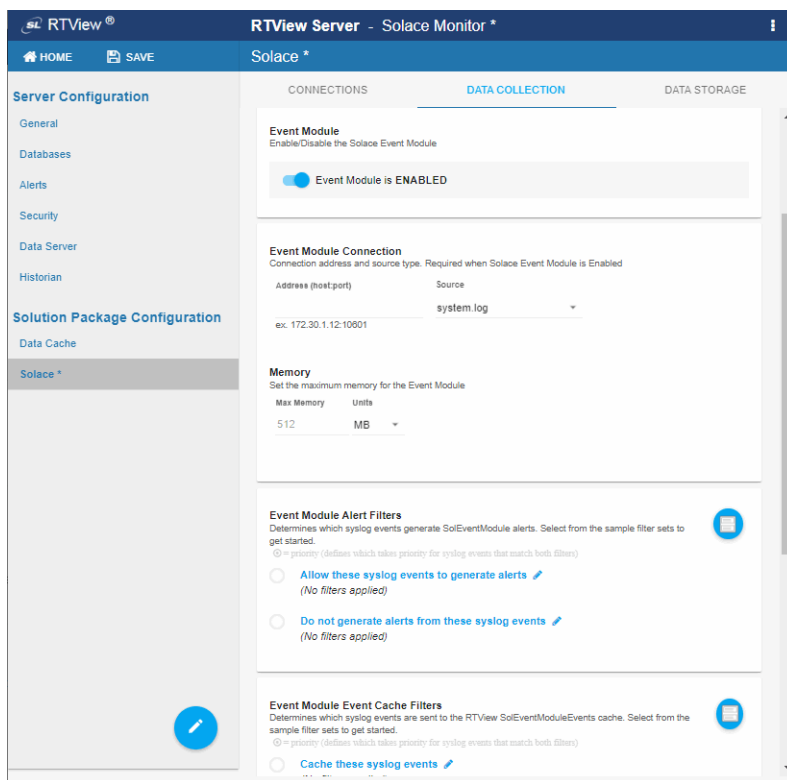
<https://docs.solace.com/System-and-Software-Maintenance/Monitoring-Events-Using-Syslog.htm>.



Configure Solace Event Module

1. Stop RTView by executing the **SolacePubSubMonitor\bin\stop_servers.bat/sh** script.
2. “[Open the RTView Configuration Application](#)”, select **Solace** (in the navigation tree) > **DATA COLLECTION** tab and toggle ON **Event Module** (toggle is blue when enabled).



3. Scroll down to see Solace Event Module options that become available.



4. Select **event.log** or **system.log** for the Source to match the settings on your broker and enter the address.
 5. Optionally enter the **Max Memory** for the Solace Event Module application and select the Units (**MB** or **GB**). Note that this must not be set to a value smaller than **512M** or to a value higher than the resources allocated for your JVM.
 6. The Solace Event Module generates alerts based on incoming Syslog events as defined by Solace. See [“Solace Event Module Caches and Alerts”](#) for more information. Under **Event Module Alert Filters** you can optionally filter which Syslog events will generate alerts.
Click  to choose from a list of sample filters or enter your own. Filters must use XPath 1.0 expressions that return a boolean value. Multiple expressions can be separated by commas. For example, you might choose to filter out alerts based on severity or scope using the **Do not generate alerts from these syslog events** filter. Exceptions to the rules for filtering out should be entered in **Allow these syslog events to generate alerts**.
 7. Solace events are stored in the SolEventModuleEvents cache. See [“Solace Event Module Caches and Alerts”](#) for more information. Under **Event Module Event Cache Filters** you can optionally filter which syslog events are sent to the SolEventModuleEvents cache. To do so:
Click  to choose from a list of sample filters or enter your own. Filters must use XPath 1.0 expressions that return a boolean value. Multiple expressions can be separated by commas. For example, you might choose to filter out events based on severity or scope using the **Do not cache these syslog events** filter. Exceptions to the rules for filtering out should be entered into the **Cache these syslog events**.
 8. Some alerts generated by the **Solace Event Module** are cleared based on another event. Others do not have a corresponding clear event. Under **Event Module Alerts** you can optionally set the amount of time to wait, in seconds, before clearing a not clearable alert. The default is **86400**.
 9. Some alerts generated by the Solace Event Module are cleared based on another event. Under **Event Module Duration** you can optionally set the amount of time, in seconds, after an event that generates a clearable alert is received to wait for the corresponding clear event before generating the alert. If the corresponding clear event is received during this time, no alert is generated. The default is **30**.
10. **Save and Restart Server** to save and apply your changes.

Solace Event Module Caches and Alerts

The Syslog events from the Solace Event Module are stored in the SolEventModuleEvents cache and displayed in the **Solution Package for Solace/Syslog Events** display.

This cache keeps history and can optionally be stored to the history database by the Historian. You can adjust the history settings in the RTView Configuration Application using the following fields in the **Solace>DATA-STORAGE** tab:

Size>History Rows – sets the maximum number of rows to keep in memory for this cache

Duration>Expire Time For Solace Event Module Events – the time (in seconds) between updates to expire rows in this cache

Duration>Delete Time For Solace Event Module Events – the time (in seconds) between updates to delete rows in this cache

History Storage>Syslog Events – toggle to true to have the Historian store this cache to the History database

The SolEventModuleAlerts Cache

The Solace Event Module generates alerts based on incoming Syslog events. The definitions of which Syslog events generate which alerts were provided by Solace. Events and alerts are scoped to SYSTEM, VPN and CLIENT. Some alerts are clearable. In this case, a syslog message generates an alert and another syslog message clears the alert.

Other alerts are not clearable. In this case, a syslog message generates the alert, but there is no corresponding syslog message to clear it. In this case, the alert automatically clears in RTView after 24 hours.

The

SolacePubSubMonitor\rtvapm\solmon\soleventmodule\config\events\event_details.json file lists all events that generate alerts. The **SolacePubSubMonitor\solmon\soleventmodule\config\events\event_correlation.json** file lists the raising and clearing events for all clearable alerts. These files must not be modified, but used for reference only.

Alert events from the Solace Event Module are stored in the SolEventModuleAlerts cache.

This cache is not visible in displays, nor does it store history. It is used solely to generate the three alerts described below.

Alert events are deleted from this cache within 2 minutes of the alert being cleared. These alert events are cleared in 2 ways:

- Clearable alerts are cleared when the clearing event is received by the Solace Event Module.
- Non-clearable alerts are cleared after 24 hours. This time can be adjusted in the Configuration Application under **Solace>DATA-COLLECTION>Solace Event Module Alerts>Clear Time**. The value is in seconds.

The following RTView alerts are generated from the SolEventModuleAlerts cache:

SolEventModuleBrokerAlert – This alert is generated for rows in the SolEventModuleAlerts where Scope=SYSTEM. In Enterprise Monitor, this alert is mapped to the SOLACE-MSGROUTER CType.

SolEventModuleVpnAlert - This alert is generated for rows in the SolEventModuleAlerts where Scope=VPN. In Enterprise Monitor, this alert is mapped to the SOLACE-VPN CType.

SolEventModuleClientAlert - This alert is generated for rows in the SolEventModuleAlerts where Scope=CLIENT. In Enterprise Monitor, this alert is mapped to the SOLACE-CLIENT CType.

When notifying on alerts using the Java Command option, you can get additional information about the SolEventModule alerts as follows:

1. Edit

projects\custom\src\com\sl\rtvapm\custom\RtvApmCommandHandler.java to uncomment the if statement that calls `getEventModuleInfo` at the end of the `outputAlertNotification` method and also uncomment the `getEventModuleInfo` method.

2. Edit `projects\custom\src\make_all.bat` or `make_all.sh` to add `RTVAPM_HOME/rtvapm/solmon/lib/rtvapm_solmon.jar` to the CP.
3. Run `make_all.bat` or `make_all.sh` to rebuild the jar.
4. In the RTView Configuration Application, configure your alert notifications to execute a Java command (which will generate a message in the data server log file for each alert notification). For SolEventModule alerts, there will be a second message containing the information from the corresponding row in the SolEventModuleAlerts cache.

Solace Event Module Logging

The Solace Event Module generates a separate log file from the data server under `logs\soleventmodule.log`. You can adjust the logging for the solace event module by modifying the `projects/rtview-server/soleventmod.log4j2.properties` file.

High Availability

High Availability (HA) mitigates single point of failure within the Solution Package for Solace system by providing a means of defining redundant system components, together with failover capability, for users of those components.

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

This section contains the following:

- [“HA Architecture”](#)
- [“Requirements”](#)
- [“Configure HA”](#)
- [“Verify HA Setup”](#)

HA Architecture

Data Server HA

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

Display Server HA (Classic UI -RTView Manager Only)

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

HTML UI HA (Solace PubSub+ Monitor UI)

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

Historian HA

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

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

Requirements

The following are minimum requirements for High Availability:

- Two host machines, one for the primary host and one for the backup host.
- Both hosts must be configured such that the RTView processes on each host can connect to each other via socket.
- Both hosts must be able to access:
 - the same data connections
 - the same historian database
 - the alert threshold database
- The RTView processes on both hosts must be able to run against identical properties files. In the case where drivers or other third party jars are located in different directories on the two hosts, create a directory in the same location in each host, copy the jar files into and reference that directory in your properties.
- Tomcat or other Application Server that can access both the primary host and backup host.

Configure HA

To configure high availability:

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

Windows:

 - From the command line on the primary host, type **bin\start_servers -haprimary**.
 - From the command line on the backup host, type **bin\start_servers -habackup**.

UNIX:

 - From the command line on the primary host, type **bin/start_servers.sh -haprimary**.
 - From the command line on the backup host, type **bin/start_servers.sh -habackup**.

Verify HA Setup

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

- ["Primary Data Server Log File"](#)
- ["Backup Data Server Log File"](#)
- ["Primary Historian Log File"](#)
- ["Backup Historian Log File"](#)
- ["Primary Display Server Log File"](#)
- ["Backup Display Server Log File"](#)

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

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

Primary Data Server Log File

```
startup
[rtview] Starting as primary HA data server accessible via //primaryhostname:4178,//
backuphostname:4178
[rtview] DataServerHA: connected to backuphostname:4178
[rtview] DataServerHA: run as primary server, backuphostname:4178 has lower priority than
this server
[rtview] leaving standby mode
```

Backup Data Server Log File

```
startup
[rtview] Starting as backup HA data server accessible via //primaryhostname:4178,//
backuphostname:4178
rtview] entering standby mode
after failover (primary data server exits)
[rtview] DataServerHA: error receiving message: java.net.SocketException: Connection
reset (primaryhostname:4178)
[rtview] DataServerHA: becoming primary server, lost connection to primary server
primaryhostname:4178
[rtview] leaving standby mode
after failback (primary data server comes back up)
[rtview] DataServerHA: resigning as primary server, got standby directive from other
server primaryhostname:4178
[rtview] connected to primaryhostname:4178
[rtview] entering standby mode
```

Primary Historian Log File

```
[rtview] Starting as primary HA historian paired with backup historian at
<backuphostname>:4122
[rtview] ServerGroup: status of member <backuphostname>:4122 : primary, priority= 1,
started=Wed Nov 14 12:56:01 PST 2018
[rtview] ServerGroup: primary server = local
[rtview] ServerGroup: becoming primary server
```

Backup Historian Log File

```
[rtview] Starting as backup HA historian paired with primary historian at
<primaryhostname>:4122
[rtview] ServerGroup: status of member <primaryhostname>:4122 : primary, priority= ,
started=Wed Nov 14 12:56:01 PST 2018
[rtview] ServerGroup: primary server = <primaryhostname>:4122
after failover (primary historian exits):
[rtview] error receiving message: java.io.EOFException (primaryhostname:4122 )
[rtview] ServerGroup: disconnected from primaryhostname:4122
[rtview] ServerGroup: primary server = local
after failback (primary historian starts back up):
[rtview] ServerGroup: status of member primaryhostname:4122 : primary, priority= 2,
started= Tue Nov 20 09:12:43 PST 2018
[rtview] ServerGroup: connected to primaryhostname:4122
```

```
[rtview] ServerGroup: primary server = primaryhostname:4122
```

Primary Display Server Log File

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

Backup Display Server Log File

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

Property Editor REST API

This section describes the Monitor REST API you can use to add, edit and delete properties on a running data server. This means that you can update connection properties without restarting the data server.

Note: Changes to Solace connections that are not to Cloud Brokers require restart. They are not applied when the server properties are updated. Changes to Cloud Broker connection properties do not require restart.

To complete these instructions you need the abbreviated name for the Monitor--also called the **PackageName**. The **PackageName** for Solution Package for Solace and the Solution Package for Solace is **solmon**. Where indicated, you:

replace **<PackageName>** with **solmon**

For example, change:

```
node main.js -action=getPropertyDescriptions -sp=<PackageName>
```

to:

```
node main.js -action=getPropertyDescriptions -sp=solmon
```

A sample node.js-based application is available in the **rtvapi/sampleapps/propeditor** directory which you can use to edit properties via the same rtvadmin servlet that is used by the RTView Configuration Application. This sample application also serves as an example of how to post to the rtvadmin servlet from your own application. For instructions about how to setup and run the sample application see the **README.txt** file in the same directory.

Two use cases are supported:

- **"Import Initial Properties & Connections into Configuration Application"**: Rather than manually entering each connection, you can use the REST API to import initial connections into the Configuration Application. You can subsequently edit those connections using the Configuration Application.
- **"Automate Connection Updates"**: Rather than using the Configuration Application to manage your connection properties, you can use the REST API to add, edit and delete connections. This is useful when you have an automated system for provisioning and want to automatically add monitoring as part of the provisioning process. These connections will not be included in the Configuration Application and will only be edited via the REST API.

Also see [“Design Notes”](#) for details about [“Supported API Actions”](#), [“Filenames”](#), [“Sample json”](#), [“Adding, Editing, Deleting JsonPrimitive Properties”](#), [“Adding and Editing JsonObject Properties”](#), [“Deleting JsonObject Properties”](#), [“Updating vs. Restarting Data Servers”](#) and [“High Availability”](#).

This section also contains:

- [“Import Initial Properties & Connections into Configuration Application”](#)
- [“Automate Connection Updates”](#)
- [“Encrypt Property Text”](#)
- [“Design Notes”](#)

Import Initial Properties & Connections into Configuration Application

Replace **<PackageName>** with the **PackageName** for the solution package you are configuring.

To Import Properties:

1. Install and start the Monitor.
2. Open a command prompt and navigate to the **rtvapm/sampleapps/propeditor** directory. Follow the instructions in the **README.txt** file to configure the node application to connect to the Monitor.
3. By default, all properties (including passwords) are sent to the rtvadmin servlet and on to the Data Server in plain text. You can optionally encrypt that text. See [“Encrypt Property Text”](#) for details.
4. Use the sample application to retrieve a list of solution packages in your data server as follows:
node main.js -action=getSPs
5. Use the sample application to get a list of available properties for your solution package as follows:
node main.js -action=getPropertyDescriptions -sp=<PackageName>
where **<PackageName>** is the abbreviated name for a solution package on the retrieved list.
6. Create a json file containing the connections and other properties you would like to add. Note that the file contents must be valid json. See [“Sample json”](#) for details about json properties.
7. Confirm that the Configuration Application is NOT in use.
8. Use the sample application to add the properties as follows:
node main.js -action=editProperties -filename=project -propstoadd=jsonfile.json

Note that the file name must be **project** in this use case. Otherwise, the properties will not be applied. See [“Adding, Editing, Deleting JsonPrimitive Properties”](#) for additional information.

9. Use the sample application to update or restart the data server. An update will apply connection properties. A restart is required to apply non-connection properties:

Node main.js -action=updatePropertiesOnServer

Or

Node main.js -action=restartServers

10. Now that the initial properties are imported you can use the RTView Configuration Application to edit your configuration.

Automate Connection Updates

Replace **<PackageName>** with the **PackageName** for the solution package you are configuring.

To Auto-update Connections:

1. Install and start the Monitor.
2. In a text editor, open **projects\rtview-server\rtvservers.dat** and add -
properties:autoconnections at the end of the dataserver line.
3. Open a command prompt and navigate to the **rtvapm/sampleapps/propeditor** directory. Follow the instructions in the **README.txt** file to configure the node application to connect to the Monitor.
4. By default, all properties (including passwords) are sent to the rtvadmin servlet and on to Data Server in plain text. You can optionally encrypt that text. See [“Encrypt Property Text”](#) for details.
5. Use the sample application to get a list of solution packages in your Data Server as follows:
node main.js -action=getSPs
6. Use the sample application to get the list of available properties for a solution package as follows:
node main.js -action=getPropertyDescriptions -sp=<PackageName>
where **<PackageName>** is the abbreviated name for a solution package on the retrieved list.
7. Create a json file containing the connections and other properties you would like to add. Note that the file contents must be valid json. See [“Design Notes”](#) below for details about json properties.
8. Use the sample application to add the properties as follows:
node main.js -action=editProperties -filename=autoconnections -propstoadd=jsonadd.json

Note that the file name must match the **-properties** command line argument that you entered in **rtvservers.dat**. See [“Filenames”](#) for more information.

9. Use the sample application to update or restart the data server. An update will apply connection properties. A restart is required to apply non-connection properties:

Node main.js -action=updatePropertiesOnServer

Or

Node main.js -action=restartServers

10. Now that the initial connections have been added, you can delete or modify those connections as follows:

node main.js -action=editProperties -filename=autoconnections -propstoadd=jsonadd.json -propstoremove=jsondelete.json -merge=true

See [“Design Notes”](#) for more information.

Note: In this scenario it is possible that the automated property updates occur at the same time as someone is editing other properties in the Configuration Application. Since all properties files are re-read when you execute the **updatePropertiesOnServer** post, the properties saved by the Configuration Application are re-read as well. The Configuration Application might say that you need to restart servers when it isn't necessary.

To encrypt property text, proceed to [“Encrypt Property Text”](#).

Encrypt Property Text

By default, properties (including passwords) are sent in plain text from the client application to the servlet. To use AES encryption on the text, do the following:

1. In the sample node.js-based application (in the **rtvapm\sampleapps\propeditor** directory), set the **cryptKey** variable to the key you want to use for the AES encryption.

The application might clip or pad this key as needed in order to generate a 16 element byte array that can be used by AES encryption.

2. In the data server's **rtvservers.dat** file, pass the value you used for **cryptKey** into the command line using the **-propkey** command line argument on the data server line.

You can either enter the key in plain text or you can scramble it using the **encode_string** command line utility.

For example, you could pass in **-propkey:propertyKeyValue**. Or you could scramble the key as follows on the command line: **encode_string propertyKeyValue**

which returns this value:

```
01343013550134901335013330134801335013500134601331013490134901353013450134801334.
```

You can then use that value on the command line instead: -

```
propkey:01343013550134901335013330134801335013500134601331013490134901353013450134801334
```

Design Notes

This section contains:

- ["Supported API Actions"](#)
- ["Filenames"](#)
- ["Sample json"](#)
- ["Adding, Editing, Deleting JsonPrimitive Properties"](#)
- ["Adding and Editing JsonObject Properties"](#)
- ["Deleting JsonObject Properties"](#)
- ["Updating vs. Restarting Data Servers"](#)
- ["High Availability"](#)

Supported API Actions

The REST API supports several actions. To get the list of actions, go to the sample application as described above and execute the following on the command line:

```
node main.js -action=getActions
```

To get the description of a single action:

```
node main.js -action=getActions -name=actionName
```

You can also execute any action that start with get in a browser as follows (where **host**, **port** and **rtvadmin** are the values you specified in the sample application):

```
http://host:port/rtvadmin/api?action=getActions&name=actionName
```

Filenames

When using the REST API to import initial properties into the Configuration Application, the filename must be **project**. This is because the Configuration Application reads and writes the project properties files and all RTView projects automatically read them. When using the REST API to automatically update properties that are not included in the Configuration application, the filename must match the **-properties** argument in the **rtvservers.dat** file and must NOT be **project**.

Sample json

You can optionally use the Configuration Application to generate sample json to get you started. Properties saved from the Configuration Application are in **projects\rtview-server\project.properties.json**.

Adding, Editing, Deleting JsonPrimitive Properties

All primitive json values must be enclosed in quotes, even boolean and number values. The top level solution package element must be included.

The following example uses **solmon** properties to illustrate. See the generating sample json properties for details about generating properties for your solution package.

Example:

```
{
  "solmon": {
    "expiretime": "10000"
  }
}
```

Adding and Editing JsonObject Properties

Solution package connections are arrays of JsonObjects. The property descriptions indicate which fields in the json object are required and which are indexes. When adding a new connection (or other JsonObject), you must include all of the required and index fields or the property will not be saved. The top level solution package element must be included.

The following example uses **solmon** properties to illustrate. See the generating sample json properties for details about generating properties for your solution package.

Example:

```
{
  "solmon": {
    "conn": [{
      "iscloudvmr": "true",
      "__name": "conn2",
      "url": "http://host2:8080/SEMP",
      "version": "7.4VMR",
      "vpnnamelist": "vpn1;vpn2"
    },
    {
      "iscloudvmr": "true",
      "__name": "conn3",
      "url": "http://host3:8080/SEMP"
    }
  ]
}
```

When adding connections to an existing file, you can either merge the new connections into the existing connection list or you can replace the whole list with the connections. This is controlled by the merge parameter. When merge is true, the indexes are used to control whether a new connection is added or an existing connection is modified.

Deleting JsonObject Properties

Solution package connections are arrays of JsonObjects. The property descriptions indicate which fields in the json object are indexes. When deleting a connection (or other JsonObject), only the index fields are required. The top level solution package element must be included.

The following example uses **solmon** properties to illustrate. See the generating sample json properties for details about generating properties for your solution package.

Example:

```

{
  "solmon": {
    "conn": [{
      "__name": "conn2"
    },
    {
      "__name": "conn3",
    }
  ]
}

```

Updating vs. Restarting Data Servers

All connection properties support updates. Once you have added, edited or deleted connections using the REST API, you can apply those changes with the `updatePropertiesOnServer` action. Restart is not required. Note that when connections are removed from your configuration, they are not immediately removed from the monitor. They stay in the caches (and display) but do not receive further updates. They will expire and be removed based on the settings in the DATA STORAGE tab of the Configuration Application. All non-connections properties are applied on restart, so they must be applied with the `restartServers` action. Restarting your servers will also cause any deleted connections to be immediately removed from the caches and displays.

High Availability

To edit properties for HA-configured servers, first follow the instructions in the **High Availability** section of this document to configure the `rtvadmin` servlet for High Availability.

Create Custom Alerts

This section describes how to define custom alerts. You can define custom alerts against any Solution Package metric. Once custom alerts are deployed they behave as built-in alerts do. That is, you can configure thresholds and duration in the **Alert Administration** display, they are listed in the **Alert Table** when they trigger, and you can ["Configure Alert Notification"](#) on them.

Custom alerts can optionally be associated with a Solution Package alert category and/or CI Type which means they are included in monitor **Alert Counts**.

Alerts must be defined and deployed in the RTViewDataServer installation directory containing the metrics that trigger them.

Defining Alerts

To define an alert, create a json file named **custom-alert-AlertName.json** under **projects/custom/alerts**, where **AlertName** is the name of the alert you will see in the monitor **Alert Administration** and **Alert Table**. The **AlertName** must not contain spaces.

The json file contains the definition for a single alert as described below. To define multiple custom alerts, create a separate json file for each alert.

```
{
  "alert": {
    "condition": ">",
    "metrics": [
      {"cache" : "EmsServerInfo",
       "metriccols" : "connectionCount;durableCount"}
    ],
    "calc" : "%metric1 + %metric2",
    "alert": "50",
    "warn": "25",
    "enabled": "true",
    "duration": "30",
    "warntext": "This has a warning: $alertIndex"
    "alerttext": "This has a critical alert: $alertIndex"
    "spname": "emsmon"
    "Package": "Ems",
    "Category": "Server",
    "citype": "EMS-SERVER",
  }
}
```

Some fields are required and others are optional. Note that even non-string fields must be specified as a json string (enclosed in quotes).

The following fields define the data that drives the alert.

- **condition**: Required. This defines which the condition that executes the alert. Valid values are >, <, >=, <= and =. For example, if the condition is >, then the alert will trigger when the metric value is greater than the threshold.
- **metrics**: Required. This is the cache metric that will drive the alert. It must be a one element array. If multiple elements are defined, only the first will be used. To see what metrics are available in your data server, go to the **Cache Table** view in the monitor.

In the monitor, go to the **Admin** tab and navigate to **Admin->Cache Table**. The top table lists the available caches. Click on a cache to see the cache contents in the lower table. Once you have found the metric against which you want to alert, note the name of the cache and the name of the column containing the metric to use in the **cache** and **metriccols** arguments below.

- **cache**: Required. The name of the cache containing the metrics to drive the alert. This must be a cache that is in the data server as the alert.
- **metriccols**: Required. The name(s) of one or more columns containing the metric data to drive the alert. If no calc is specified, only the first column is used. This must not contain the index columns. The columns specified here must contain numeric data for all conditions except =, in which case, it can be a number or string.
- **calc**: Optional. If specified, this calculation will be applied to the metric(s) before they are passed into the alert. The value must be an expression suitable to pass into the Evaluate Expression By Row (JEval) function. The variable names to use in the expression must be %metricX where X is the 1 for the first column in the **metrics:metriccols**, 2 is the second column in **metrics:metriccols**, etc. No validation of the expression syntax is performed so you must enter a valid value or check the log file after the alert is loaded.
- **alert**: Required. It is the default alert threshold value. This value is used when this alert is initially inserted into the ALERTLEVEL database. It must be a double for all conditions except =. If the condition is = and the **metriccols** column contains a string values, the alert will do a string comparison against this value and will trigger a severity 2 (critical) alert when the condition is met. You will not be able to set thresholds in the **Alert Admin** for string comparison alerts.
- **warn**: Required. This must be a double. It is the default warning threshold value. This value is used when this alert is initially inserted into the ALERTLEVEL database. NOTE: This is not required in the case where condition is = and alert is a string.
- **enabled**: Optional. This must be **true/false**. It is the default **Enabled** value. If not specified, **false** will be used.
- **duration**: Optional. This must be an integer. It is the default **Duration** value in seconds. If not specified, **30** will be used.
- **alerttext**: Optional. The text to use in the **Alert Text** field of the **Alert Table** for critical alerts. It can include the following substitutions to show alert info in the text: \$alertName, \$alertText, \$alertID, \$alertSeverity, \$alertIndex, \$alertLabel, \$alertCurValue, \$alertCompValue, \$alertTime, \$alertCommandText, \$alertClearedReason, \$alertClearedTime.

If not specified, the following text is shown in the **Alert Text** for critical alerts: “**Alert: 1234 value 10 > 5**”, where **1234** is the alert id, **10** is the alert value, **>** is the condition and **5** is the threshold value.

- **warntext**: Optional. The text to use in the **Alert Text** field of the **Alert Table** for warning alerts. It can use the same subs as the **alerttext**. If not specified, the following text is shown in the **Alert Text** field for warning alerts: **Warning: 1234 value 10 > 5**, where **1234** is the alert id, **10** is the alert value, **>** is the condition and **5** is the threshold value.
- **spname**: Optional. If specified, this alert will be associated with the specified Solution Package in the monitor. It will be included in the alert count in the **Overview** page for the specified Solution Package and will be filtered by the **Packages** selection in the **Alert Administration** page. It is also used in conjunction with the **Package** and **citype** fields described below. A good way to determine what value to use here is to look at the **spname** value used for the built-in alerts. To see this in the Monitor, go to the **Admin** tab under **Admin->Cache Table**. Select the RTViewDataServer for this alert and look at the RtvAlertdefMetaData cache.
- **Package**: Optional. If specified, this value sets the **Package** property on the alert. The **Package**, **Category** and **spname** values control whether or not the alert will be included in the alert count in the **Component** displays. The alert count shows up in the top right corner and counts the alerts associated with the data in that display.

For example, if the **spname** is set to **emsmon**, **Package** is set to **Ems** and **Category** is set to **Server**, this alert should be included in the alert counts for the **EMS Server Summary** pages. To see what **Package** and **Category** names are used in a given display, click on the **Alert Count**. This will drill down to an alert view that shows the **Package** and **Category**. You can also look at the values used for the built in alerts in the RtvAlertdefMetaData cache as described in the **spname** description.

- **Category**: Optional. If specified, this value sets the **Category** property on the alert. See the **Package** field description for how this is used.
- **citype**: Optional. If specified and this alert is used with RTView Enterprise, this alert will be mapped to the specified **CITYPE** in the service model and included in the alert counts for services containing the specified CITYPE. The source of data must have the same index column(s) as the CITYPE definition. To see the index columns for your alert, deploy the alert, look at the **Index Column Names** column in the RtvAlertdefMetaData cache in the row for your custom alert. To see the index column(s) for the CITYPE definitions, go to the **Components** tab **Other->System->CI Type Definitions** and look at the **INDEXNAMES** column in the **CI Type Definitions** table. If your alert index columns do not match the CITYPE index columns, you cannot map your alert to that CITYPE.
- **indexTypes**: Optional. If specified, this defines the overrides that overrides that are available in the **Alert Administration Overrides** page. Overrides are useful when the alert has multiple index columns and you want to allow the user to override the thresholds based on a subset of the indexes. The string must be in the format **OverrideLabel:indexColumn1;OverrideLabel2:indexColumn1,indexColumns2**.

To see the index columns for your alert, deploy the alert, look at the **Index Column Names** column in the RtvAlertdefMetaData cache in the row for your custom alert. For example, here is one from an EMS queue alert:

PerQueue:name;PerServerQueue:URL,name.

- **description**: Optional. A description of the alert. This value is shown in the **Alert Administration** page when the alert is selected.

Deploying Alerts

To deploy your custom alerts, restart your data server. It will automatically load all custom alert files under **projects/custom/alerts** and add a row to the ALERTLEVELS database for each custom alert using the default alert, warning, enabled and duration values specified in the json file. To remove a custom alert, remove the json file for that alert from **projects/custom/alerts** and restart your data server. This will remove the alert from the system, but will not delete it from the ALERTLEVELS database and therefore it will still show up in the Alert Administration page until you delete it from the database. To modify your custom alert, change the .json file for that alert, then restart the data server.

To deploy your custom alert to another RTViewDataServer installation, copy the **projects/custom/alerts** directory to that installation. For high-availability RTViewDataServer installations, create, deploy and validate the alerts in the primary installation, then copy the **projects/custom/alerts** directory to the backup installation. Any changes to the alerts definitions should be made on the primary data server, then copied to the backup server.

Verifying Alerts

There are 4 places to verify that your alert was configured and deployed as expected.

1. Check the log file for your data server for errors about your custom alert(s). If there was an error in the .json file that prevented the alert from being loaded, you will see a message like this: "ERROR: cannot load alert AlertName." followed by a description of the error. If the error message is "json content is invalid", this means you have a json syntax error (for example, you're missing a quote or a comma". Use an online json validator to correct your syntax.
2. Check the RtvCustomAlerts and RtvAlertdefMetaData caches for your alert(s). The RtvCustomAlerts cache shows a row for each loaded custom alert with its Package and **CITYPE** settings. The RtvAlertdefMetaData cache shows one row for all alerts in the data server, including all of the built-in alerts. Filter the table on **IsCustom = true** to show only custom alerts. It shows the full alert configuration.
3. Look for your alert in the **Alert Administration** display. It should use the default alert, warning, enabled and duration value from your alert definition. Note that if you deploy then delete a custom alert, it will not automatically be removed from the **Alert Administration** display. To remove it from **Alert Administration**, remove all rows in the ALERTLEVELS table for that alert.
4. Set the thresholds for your alert in **Alert Administration** such that an alert should trigger, then go to the **Alert** tab to view the alert. Note that once you change the alert threshold, enabled or duration values from **Alert Administration** changing the default values in the alert definition will NOT override the values set in **Alert Administration**. The default values are only used when the alert is loaded into the system the first time.

CHAPTER 5 Configure Alert Notification

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

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

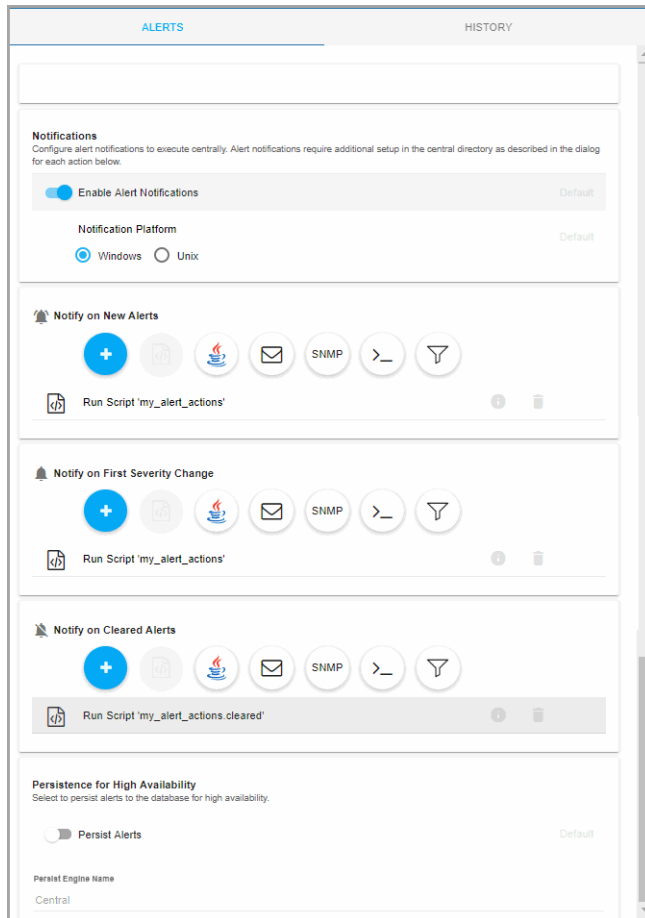
- when a new alert is created
- the first time the **Severity** level on an alert changes
- when an alert is cleared
- periodically renotify unacknowledged alerts


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

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

To configure Alert Notification:

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



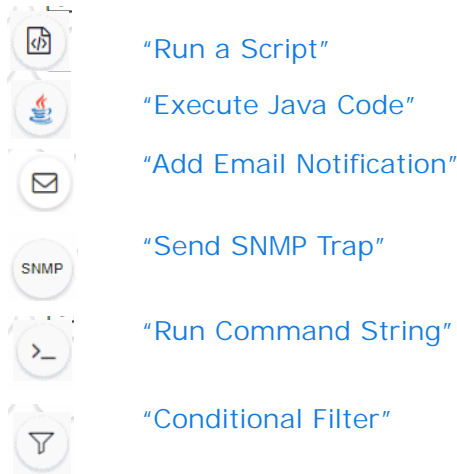
2. Toggle on **Enable Alert Notifications** and select the **Notification Platform** type (**Windows or Unix**).
3. Select an alert event that you want to notify on by clicking  next to the option.

Alert Event Options

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

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

Alert Action Options



You can choose multiple actions.

- Click **SAVE** to close the dialog and **SAVE** (in title bar) to save your changes.
- Some alert notification actions require additional setup as described in the dialog for each action. See the descriptions of each action below for details on the dialogs and additional setup for each action.
- Click **RESTART SERVERS** to apply changes.

Run a Script

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

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

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

Return to [“Alert Event Options”](#).

Execute Java Code

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

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

Make the following entries:

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

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

Customizing the Custom Command Handler

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

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

Return to ["Alert Event Options"](#).

Add Email Notification

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

Make the following entries:

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

Return to [“Alert Event Options”](#).

Send SNMP Trap

This alert notification action sends an SNMP Trap as described in **rtvapm/rtview/lib/SL-RTVIEW-MIB.txt**. This action can be added multiple times per notification type. No additional setup is required beyond filling in the **Add SNMP Trap Notification** dialog in the Configuration Application.

Make the following entries:

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

Return to [“Alert Event Options”](#).

Run Command String

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

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

Return to [“Alert Event Options”](#).

Conditional Filter

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

To create a condition, make the following entries:

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

Return to ["Alert Event Options"](#).

CHAPTER 6 Using the Monitor

The Solution Package for Solace is an advanced messaging platform that allows customer applications to efficiently exchange messages over dedicated VPNs. The Solution Package for Solace provides pre-configured alerts and dashboards to monitor current status and manage history for the Solace broker. The Solution Package 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 contains:

- [“Login to Solution Package for Solace”](#): Describes how to access the Solution Package for Solace and [“User Permissions”](#).
- [“Overview”](#): Describes the Monitor [“Graphic Elements”](#) and functionality.
- [“Displays”](#): Describes Monitor displays under the [“Displays”](#) tab.
- [“Alerts”](#): Describes Monitor displays under the [“Alerts”](#) tab.
- [“Admin”](#): Describes Monitor displays under the [“Admin”](#) tab.

Login to Solution Package for Solace

To access Solution Package for Solace to monitor your Solace components, browse to:

http://IPAddress:8068/rtview-solmon if you are executing your browser on a different host than where the monitor is running.

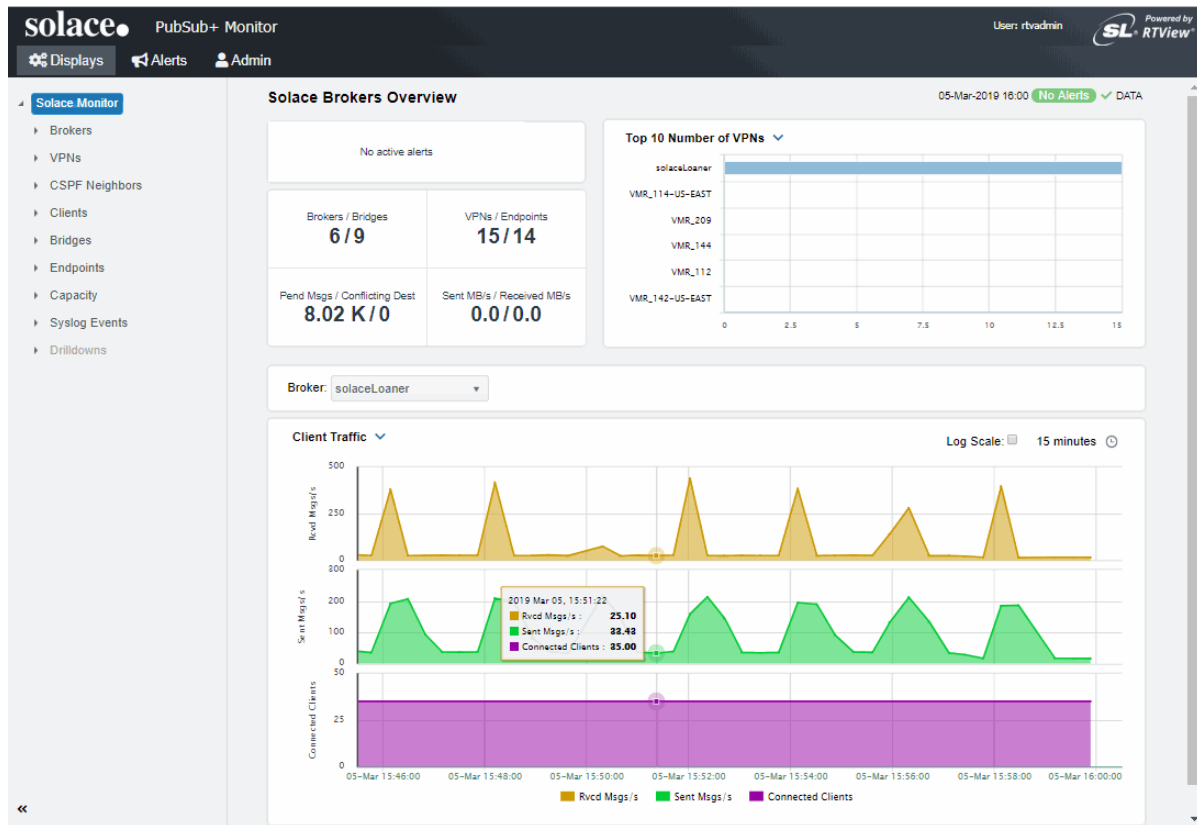
http://localhost:8068/rtview-solmon if you are executing your browser in the same host where the monitor is running.

User Permissions

There are three types of users:

- **End-users** use `rtvuser/rtvuser` as their username/password which permits read-only access to all displays except for **Admin** tab displays.
- **End-user with alert management privileges** use `rtvalertmgr/rtvalertmgr` as their username/password which permits the same access as the end-user. Additionally, you can use the **Own**, **Ack**, **Unack** and **Comment** functions in the **Alerts Table**.
- **Administrators** use `rtvadmin/rtvadmin` as their username/password which permits read-only access to all displays as well as **Admin** tab displays. You can also enable and administer alerts, view cache contents and use the **Own**, **Ack**, **Unack** and **Comment** functions in the **Alerts Table**.

The Solace PubSub+ Monitor Displays home page opens, which provides a health summary of all your Solace brokers (see the following figure).



On larger screens the page contains a horizontal menu bar with three tabs:

- **Displays** contains the screens for PubSub+ performance data which you select from the navigation tree in the left panel.
- **Alerts** is used for viewing and managing alerts.
- **Admin** is used for administering alerts and viewing cache contents directly. This tab is only accessible to users with administrator privileges (user accounts with the rtvadmin role). You can hide the navigation tree by clicking << (on the lower left).

Navigation through the displays is recorded in the browser history and you can use the browser's back and next buttons to traverse that history. You can hide the navigation tree in the **Displays** and **Admin** tabs by clicking << (on the lower left).

On smaller screens, the horizontal menu bar is replaced by a vertical menu whose visibility is toggled by clicking the menu icon in the upper right corner of the page.

Once a user is logged in, that user remains logged in until the browser window is closed. Closing just the browser tab that contains the user interface does not log out the user, the browser itself must be closed.

See ["Displays"](#) for details about displays for Solace PubSub+.

Overview

This section describes the general operation of the Solution Package for Solace, the user interface as well as “Graphic Elements” such as “Heatmaps”, “Tables”, “Trend Graphs” and “Icons and Buttons”.


Graphic Elements


This section describes the graphic objects that are used in displays and their behavior:

- “Heatmaps”
- “Tables”
- “Trend Graphs”
- “Icons and Buttons”

Heatmaps

Heatmaps organize your Solace PubSub+ resources (brokers, VPNs, Clients, Bridges, Endpoints and so forth) 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 Brokers Heatmap** is shown below).



For example, the **Solace Brokers 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 a Solution Package for Solace 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 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.

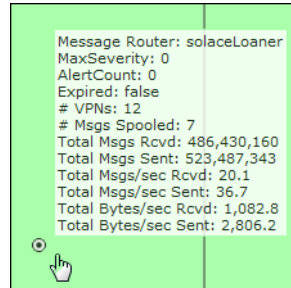
● Metrics that have not exceeded either specified threshold have an **Alert Severity** value of **0**. For a given rectangle, this indicates that no metrics have reached their warning or alert thresholds.

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.


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

Solution Package for Solace 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 **Solace Brokers Table** display (shown below) shows the same data as the **Solace Brokers Heatmap** display. The following figure also illustrates the "Column Visibility" which allows you to select the columns you want in the table.



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"](#)
- ["Column Visibility"](#)
- ["Column Filtering"](#)
- ["Column Reordering"](#)
- ["Row Paging"](#)

Multiple Column Sorting

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


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.



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.

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



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.

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.

Row Paging


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

217	emreference	sl.rtvew.sql.sqldb	sl.rtvew.sub	\$rtvConfigDataServer.CONFIG_SERVER
229	emreference	sl.rtvew.properties.queryTimeOut	10	
216	emreference	sl.rtvew.sql.sqldb	ALERTDEFS --- __none ---	

Page 1 of 2 1 - 200 of 235 items

Trend Graphs


Solution Package for Solace 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.

Icons and Buttons

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.

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 filter the item/s you want to view. Options differ among displays.

Displays

The Solace PubSub+ Monitor organizes displays under the following Views:

- **"Brokers"**: The displays in this View present broker-level metrics, which reflect configuration settings, total throughput, current status, errors, and value-added calculations that summarize metrics across all of the VPNs.
- **"CSPF Neighbors"**: The displays in this View present a topology and metrics of your brokers, PubSub+ 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 broker. 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 broker, 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 broker level.
- **"Syslog Events"**: View details about Syslog events.
- **"Drill Down Displays"**: These displays are accessed via other displays (with the exception of **"Alerts History Table - HTML"**).

Brokers

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

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

Brokers Overview

The **Brokers Overview** is the top-level display, which provides a good starting point for immediately getting the status of all your brokers on your Data Server.

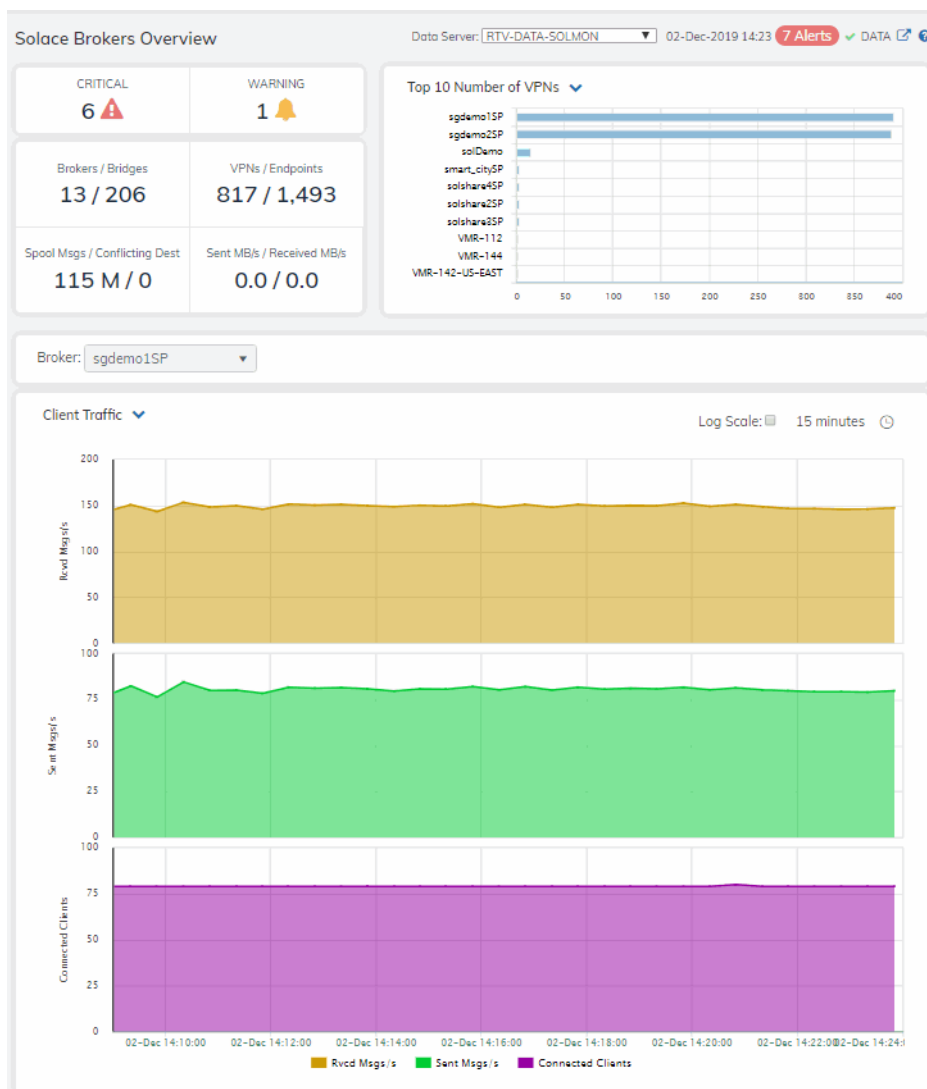
Select a data server, broker and metric from the drop-down menus. Consider keeping this display open for monitoring at a glance. You can easily view the current data for that Data Server including:

- **Top 10** most utilized **VPNs / Endpoints**, **Clients Connected** and **Spooled Messages**.
- The number of **Brokers / Bridges**.
- The number of **Spooled Messages / Conflicting Destinations**.
- The number of **Sent MBs per second / Received MB per second**.

You can hover over each area in the upper half of the Overview to see more detail. You can also drill down to see even more detail by clicking on each metric card in the Overview.

The bottom half of the display provides a performance trend graph for queries for a selected broker. The trend graph traces the performance metric you select: **Client Traffic**, **Spool Msgs** or **Memory**.

You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.




CRITICAL	Total number of current critical alerts for brokers on the selected data server.
WARNING	Total number of current critical alerts for brokers on the selected data server.
Brokers/Bridges	Total number of brokers/bridges on the selected data server.
VPNs/Endpoints	Total number of VPNs/endpoints on the selected data server.
Spooled Msgs/Conflicting Dest	Total number of spooled messages/conflicting destinations on the selected data server.
Sent MBs/Received MBs	Total number of MBs sent/MBs received on the selected data server.
Top 10 Number of VPNs	Ten brokers with the greatest number of connected VPNs.
Broker	Select a broker to trace performance metrics in the trend graph, then choose a metric:



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


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



Time Settings

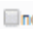


By default, the time range end point is the current time. To change the time range, click the **Time Settings**  and either:

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



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


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

Log Scale

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




Brokers Heatmap

View the current status and alerts in a heatmap of all brokers or a subset of brokers. Use the **Show** dropdown menu to choose **All** brokers, **Expired** brokers, **Unexpired** brokers or only brokers in **Standby** mode

Each rectangle in the heatmap is a single broker where the rectangle size represents the number of connections. The rectangle color maps where the current value is on its color gradient  bar. Select a broker from the drop-down menu. For example, by default, **Alert Severity** is shown:

Alert Severity


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













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

Click a rectangle to drill down to details about a broker in the [“Broker Summary”](#) display.

Mouse over a rectangle to see additional details. Use the check-box ☒ to include / exclude **Connected** brokers and enable **Log Scale** mode.

Consider keeping this display open for monitoring your Solace brokers at a glance.



Alert Severity	<p>The current alert severity. Values range from 0 - 2, as indicated in the color gradient  bar, where 2 is the highest Alert Severity:</p> <ul style="list-style-type: none">  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	<p>The total number of critical and warning alerts. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of alerts in the heatmap. The middle value in the gradient bar indicates the average alert count.</p>
# Msgs Spooled	<p>The total number of spooled messages. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of SolMsgRouterPendingMsgsHigh. The middle value in the gradient bar indicates the middle value of the range.</p>
Total Msgs Rcvd	<p>The total number of received messages. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of total messages received in the heatmap. The middle value in the gradient bar indicates the average count.</p>
Total Msgs Sent	<p>The total number of sent messages. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the maximum count of total messages sent in the heatmap. The middle value in the gradient bar indicates the average count.</p>
Total Msgs/ sec Rcvd	<p>The number of messages received per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of SolMsgRouterInboundMsgRateHigh. The middle value in the gradient bar indicates the middle value of the range.</p>
Total Msgs/ sec Sent	<p>The total number of messages sent per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of SolMsgRouterOutboundMsgRateHigh. The middle value in the gradient bar indicates the middle value of the range.</p>
Total Bytes/ sec Rcvd	<p>The total number of bytes received per second in the broker. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of SolMsgRouterInboundByteRateHigh. The middle value in the gradient bar indicates the middle value of the range.</p>
Total Bytes/ sec Sent	<p>The total number of bytes sent per second in the broker. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of SolMsgRouterOutboundByteRateHigh. The middle value in the gradient bar indicates the middle value of the range.</p>

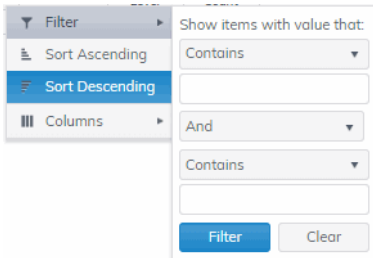
Brokers Table

Investigate detailed utilization metrics for all brokers. This display provides a tabular view of the performance metrics shown in the [“Brokers Heatmap”](#) (alert level, alert count, and so forth), but with additional metrics such as **Egress** and **Ingress** values.

Use the **Show:** dropdown menu to view the current status of **All** brokers, **Expired** brokers, **Unexpired** brokers or just brokers in **Standby** mode.

Each row in the table contains data for a particular broker. Click a column header to sort column data in ascending or descending order. Double-click on a table row to drill down to the [“Broker Summary”](#) display and view metrics for that particular broker. Toggle between the commonly accessed Table and Heatmap displays by clicking the drop down list on the display title.

Search by clicking the right side of a column heading/**Filter** to open the Search, Sort and Choose Columns dialog:



The dialog box is titled 'Filter' and contains the following elements:

- A dropdown menu labeled 'Filter' with a right-pointing arrow.
- A section titled 'Show items with value that:' containing two dropdown menus, both set to 'Contains', and two empty text input fields.
- A section titled 'And' with a dropdown menu set to 'And' and an empty text input field.
- Buttons for 'Filter' and 'Clear'.
- On the left side, there are three tabs: 'Filter' (selected), 'Sort Ascending', 'Sort Descending', and 'Columns'.

Brokers: (in the upper right portion) is the number of brokers in the display.

Use the check-boxes ☒ to include / exclude **Connected** and **Expired** brokers.

Export to Excel by right-clicking a column heading.

Toggle between **More Columns** / **Fewer Columns**

[More Columns](#)






Solace Brokers Table 11-Jul-2019 11:57 1 Alert DATA 🔗

Show Connected Only: ☐ Show Expired: All Show Standby: All Brokers: 23

[Fewer Columns](#)

Broker	Connected	Alert Level	Alert Count	Host Name	Host Address	VPNs	Total Clients	Total Clien Connecti
Sol_Mule_Azure			0	solhcdemo0.francecen		0	0	
Sol_Mule_GCP			0	35.234.63.231	35.234.63.231	0	0	
Sol_Mule_SC			0	vmr-mr8v6yiwiawhm	34.227.76.129	1	0	
solDemo			0	solace	192.168.220.5	15	40	
Team-2-Ali-cloud			0	47.74.235.254	47.74.235.254	0	0	
Team-2-AWS			0	ec2-35-177-122-45.e	35.177.122.45	0	0	
Team-2-Azure			0	demotion-team20.sov		0	0	
Team-2-Google-cloud			0	35.187.64.112	35.187.64.112	0	0	
Team-2-lab-appliance			0	london.solace.com	217.196.247.77	1	0	
Team-2-Solace-cloud			0	mr-xy4p45t57.messag		0	0	
Team-4-PCF-demo			0	shared-vmr-1.system	35.201.65.176	0	0	
Team-6-Alibaba			0	47.74.235.254	47.74.235.254	0	0	
Team-6-AnalyticsVMR			0	54.179.163.185	54.179.163.185	0	0	
Team-6-OnPremVMR-backup			0	54.191.207.187	54.191.207.187	0	0	
Team-6-OnPremVMR-primar			0	34.214.62.219	34.214.62.219	0	0	
Team-6-PCF			0	shared-vmr-1.system	35.201.65.176	0	0	
Team-6-SolaceCloud			0	mr-91b692durd.mess		0	0	
VMR-112			0	ip-172-30-1-112	172.30.1.112	2	11	

Column Values

Broker	The name of the broker.
Connected	<p>The broker state:</p> <ul style="list-style-type: none">  Red indicates that the broker is NOT connected.  Green indicates that the broker is connected.
Alert Severity	<p>The current alert severity:</p> <ul style="list-style-type: none">  Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.  Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.  Green indicates that no metrics have exceeded their alert thresholds..
Alert Count	The total number of alerts.
Expired	When checked, performance data about the sensor has not been received within the time specified.
Host Name	The name of the host.
Platform	The name of the platform.
OS Version	The version of the operating system.
Up Time	The amount of time that the broker has been up and running.
VPNs	The total number of VPNs configured on the broker.
Total Clients	The total number of clients associated with the broker.
Total Clients Connected	The total number of clients that are currently connected to the broker.
Clients Using Compression	The number of clients who send/receive compressed messages.
Clients Using SSL	The number of clients using SSL for encrypted communications.
Max Client Connections	The maximum number of available client connections.
Endpoints	The total number of endpoints configured on the broker.
Bridges	The total number of bridges configured on the broker.
Local Bridges	The total number of local bridges configured on the broker.
Remote Bridges	The total number of remote bridges configured on the broker.
Remote Bridge Subscriptions	The total number of remote bridge subscriptions configured on the broker.
Routing Enabled	This check box is checked when the broker is configured to route messages to other brokers.
Routing Interface	The name of the interface configured to support message routing.

Total # Conflicting Destinations	The total number conflicting destinations.
SpooledSpooledMessages	The number of spooled messages on the broker.
Total Client Msgs Rcvd	The total number of client messages received on the broker.
Total Client Msgs Sent	The total number of client messages sent by the broker.
Total Client Msgs Rcvd/sec	The total number of client messages received per second by the broker.
Total Client Msgs Sent/ sec	The total number of client messages sent by the broker.
Total Client Bytes Rcvd	The total number of client bytes received by the broker.
Total Client Bytes Sent	The total number of client bytes sent by the broker.
Total Client Bytes Rcvd/sec	The total number of client bytes received per second by the broker.
Total Client Bytes Sent/sec	The total number of client bytes sent per second by the broker.
Total Client Direct Msgs Rcvd	The total number of direct client messages received by the broker.
Total Client Direct Msgs Sent	The total number of direct client messages sent from the broker.
Total Client Direct Msgs Rcvd/sec	The total number of direct client messages received per second by the broker.
Total Client Direct Msgs Sent/sec	The total number of direct client messages sent per second by the broker.
Total Client Direct Bytes Rcvd	The total number of direct client bytes received by the broker.
Total Client Direct Bytes Sent	The total number of direct client bytes sent by the broker.
Total Client Direct Bytes Rcvd/sec	The total number of direct client bytes received per second by the broker.
Total Client Direct Bytes Sent/sec	The total number of direct client bytes sent per second by the broker.
Total Client Non-Persistent Msgs Rcvd	The total number of non-persistent client messages received by the broker.
Total Client Non-Persistent Msgs Sent	The total number of non-persistent client messages sent by the broker.
Total Client Non-Persistent Msgs Rcvd/sec	The total number of non-persistent client messages received per second by the broker.
Total Client Non-Persistent Msgs Sent/ sec	The total number of non-persistent client messages sent per second by the broker.
Total Client Non-Persistent Bytes Rcvd	The total number of non-persistent client bytes received by the broker.
Total Client Non-Persistent Bytes Sent	The total number of non-persistent client bytes sent by the broker.
Total Client Non-Persistent Bytes Rcvd/sec	The total number of non-persistent client bytes received per second by the broker.
Total Client Non-Persistent Bytes Sent/sec	The total number of non-persistent client bytes sent per second by the broker.

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

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

Swap Free (KB)	The total amount of available swap (in kilobytes) on the broker.
Swap Used (KB)	The total amount of swap used (in kilobytes) on the broker.
Swap Used %	The percentage of total available swap that is currently being used.
Subscription Mem Total (KB)	The total amount of available memory (in kilobytes) that can be used by queue/topic subscriptions.
Subscription Mem Free (KB)	The current amount of available memory (in kilobytes) that can be used by queue/topic subscriptions.
Subscription Mem Used (KB)	The current amount of memory (in kilobytes) being used by queue/topic subscriptions.
Subscription Mem Used %	The percentage of available memory being used by queue/topic subscriptions.
Chassis Product Number	The product number of the chassis in which the broker is contained.
Chassis Revision	The revision number of the chassis.
Chassis Serial	The serial number of the chassis.
BIOS Version	The basic input/output system used by the chassis.
CPU-1	The name of the central processing unit (CPU 1) used by the broker.
CPU-2	The name of the central processing unit (CPU 2) used by the broker.
Operational Power Supplies	The number of available power supplies that are operational on the chassis.
Power Redundancy Config	The configuration used by the backup broker.
Max # Bridges	The maximum number of bridges allowed on the broker.
Max # Local Bridges	The maximum number of local bridges allowed on the broker.
Max # Remote Bridges	The maximum number of remote bridges allowed on the broker.
Max # Remote Bridge Subscriptions	The maximum number of remote bridge subscriptions allowed on the broker.
Redundancy Config Status	The status of the redundancy configuration.
Redundancy Status	The status of the redundant broker.
Redundancy Mode	Refer to Solace documentation for more information.
Auto-revert	Refer to Solace documentation for more information.
Mate Router Name	If redundancy is configured, this field lists the redundant broker name (mate broker name).
ADB Link Up	This check box is checked if a broker is set up to use guaranteed messaging and an Assured Delivery Blade (ADB) is set up and working correctly.

ADB Hello Up	Refer to Solace documentation for more information.
Pair Primary Status	The primary status of the broker and its redundant (failover) mate.
Pair Backup Status	Refer to Solace documentation for more information.
CPU Cores	This value is retrieved by the "show system" SEMP request. Refer to Solace documentation for more information.
Expired	When checked, performance data about the broker has not been received within the time specified.
Time Stamp	The date and time the row of data was last updated.

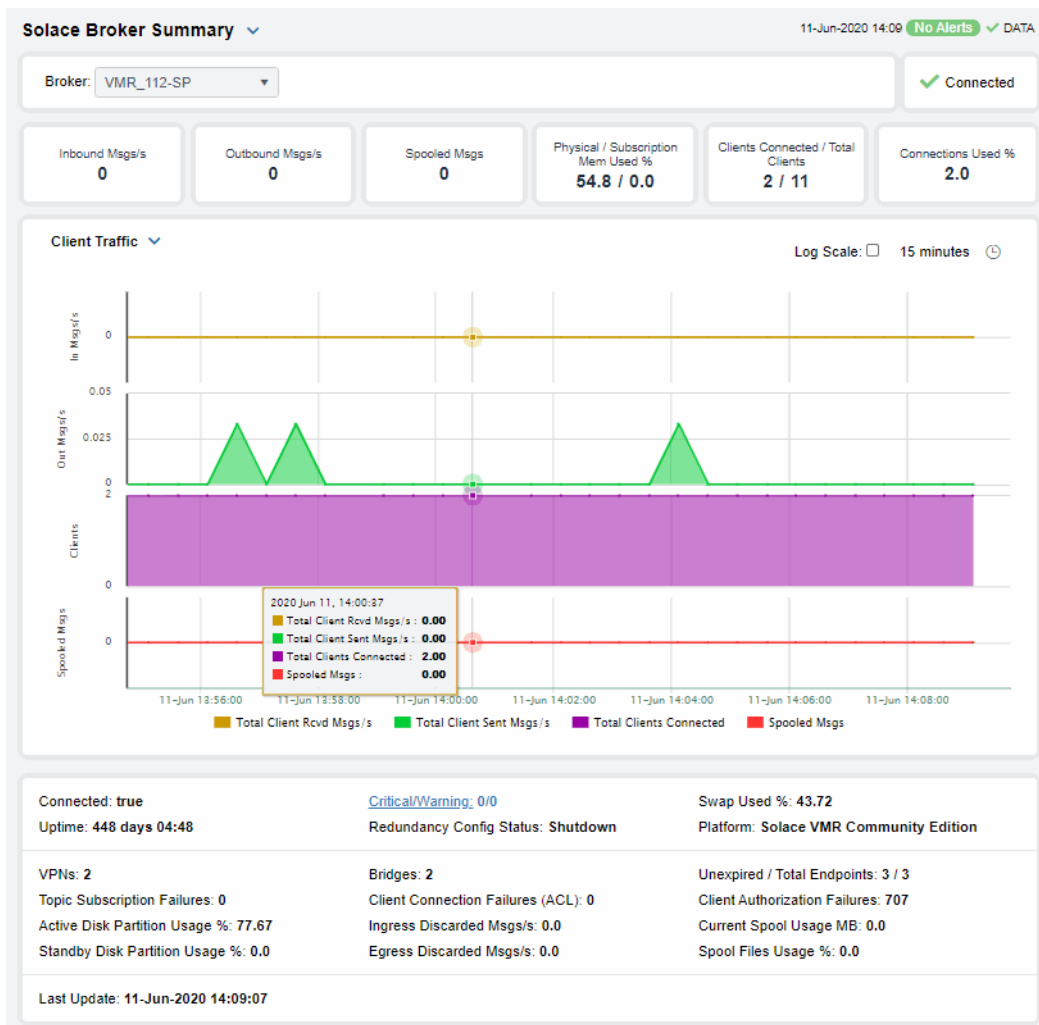
Broker Summary

View performance and processing details for a single broker, such as the total **Inbound / Outbound Messages per second**, **Spooled Messages** and **Clients Connected / Total Clients**. The trend graph traces the performance metric you select: **Client Traffic**, **Spool Msgs** or **Memory**.

Choose a broker from the **Broker** drop-down menu to view its total number of connected clients, number of incoming messages, **Up Time**, and additional information. You can also view alert statuses and **Spool Status** data for the broker. You can hover over each area in the upper half of the display to see more detail. You can also drill down to see even more detail by clicking on each metric card.

The bottom half of the display provides current and historical performance metrics for the selected broker, such as connection status and **Unexpired / Total Endpoints**.

You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.



The connection status (connected/disconnected).



Inbound Msgs/s

The number of messages received per second.

Outbound Msgs/s

The number of messages sent per second.

Spooled Msgs/s

The number of spooled messages.

Physical / Subscription Mem Used %

The total percentage of physical memory used / the total percentage of subscription memory used.

Clients Connected / Total Clients

The current number of clients connected / the total number of clients.

Connections Used %

The percentage of connections used.

Trend Graphs

Traces the selected broker.

Client Traffic

- **In Msgs/s** - Traces the total number of client messages received per second.
- **Out Msgs/s** - Traces the total number of client messages sent per second.
- **Clients** - Traces the total number of connected clients.
- **Spooled Msgs** - Traces the total number of spooled messages.

Spool Msgs

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

Memory

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



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


Log Scale

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

Time Settings


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

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



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



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

Broker Sensors

This tabular display contains environmental sensor metrics for a selected broker. Use this display to find out the type, name, value, and status of the sensors. This display only applies to Hardware (HW) Brokers. Note that the drop down menu does not show connection strings to PubSub+ Software Brokers.

Select a HW broker from the drop-down menu. Search by clicking the right side of a column heading/**Filter** to open the Search, Sort and Choose Columns dialog:

The dialog shows the 'Filter' tab selected. It includes a search bar with the text 'Show items with value that:', two dropdown menus both set to 'Contains', and two input fields for search criteria. At the bottom are 'Filter' and 'Clear' buttons.

Solace Broker Environmental Sensors 04-Mar-2019 14:38 Alerts DATA

Broker: solaceLoaner

Sensor Readings

Type	Sensor Name	Value	Units	Status	Expired	Time Stamp
Voltage	BB +1.5V	1.469	volts	OK		04-Mar-2019 14:37:52
Voltage	BB +1.5V AUX	1.490	volts	OK		04-Mar-2019 14:37:52
Voltage	BB +1.5V ESB	1.482	volts	OK		04-Mar-2019 14:37:52
Voltage	BB +1.8V	1.803	volts	OK		04-Mar-2019 14:37:52
Voltage	BB +12V AUX	12.090	volts	OK		04-Mar-2019 14:37:52
Voltage	BB +3.3V	3.337	volts	OK		04-Mar-2019 14:37:52
Voltage	BB +3.3V STB	3.337	volts	OK		04-Mar-2019 14:37:52
Voltage	BB +5V	5.070	volts	OK		04-Mar-2019 14:37:52
ThermalMargin	CPU1 Therm Margin	-87.000	degrees C			04-Mar-2019 14:37:52
ThermalMargin	CPU2 Therm Margin	-59.000	degrees C			04-Mar-2019 14:37:52
Temperature	Chassis Temp.	23.000	degrees C			04-Mar-2019 14:37:52
Fan speed	Chassis Fan 1	7714	RPM			04-Mar-2019 14:37:52
Fan speed	Chassis Fan 2	8057	RPM			04-Mar-2019 14:37:52
Fan speed	Chassis Fan 3	7714	RPM			04-Mar-2019 14:37:52
Fan speed	Chassis Fan 4	7543	RPM			04-Mar-2019 14:37:52
Fan speed	Chassis Fan 5	7371	RPM			04-Mar-2019 14:37:52
Fan speed	Chassis Fan 6	7288	RPM			04-Mar-2019 14:37:52
Power system status	Power Redundancy	yes				04-Mar-2019 14:37:52

Sensor Readings

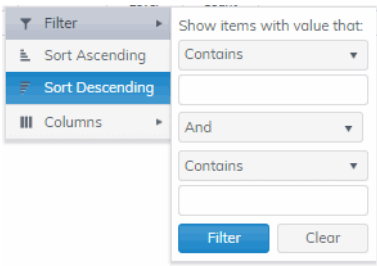
Each row in the table is a different sensor on the broker.

Type	See vendor documentation for details.
Sensor Name	The name of the sensor.
Value	Lists the value of the sensor.
Units	Lists the unit of measure for the sensor.
Status	The current status of the sensor.
Expired	When checked, performance data about the broker has not been received within the time specified.
Time Stamp	The date and time the row of data was last updated.

Broker Provisioning

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

Select a broker from the drop-down menus. Search by clicking the right side of a column heading/**Filter** to open the **Search, Sort and Choose Columns** dialog:



Broker Provisioning

04-Mar-2019 14:39 No Alerts DATA

Broker: VMR_112-SP

Host Name: ip-172-30-1-112

CPU-1: Not Specified

CPU-2: Not Specified

Platform: Solace VMR Community Edition

OS Version: soltr_7.2.1.617

Chassis Product Number: Xen - HVM domU

BIOS Version: 4.2.amazon

Chassis Revision: Unknown

Chassis Serial: Unknown

Total Memory (KB): 4,038,412

Memory Used (KB): 3,922,688

Memory Used %: 54.76

Swap (KB): 2,047,996

Swap Used (KB): 895,468

Swap Used %: 43.72

Operational Power Supplies: 0

Mate Router Name:

Power Redundancy Config: N/A

Redundancy Config Status: Shutdown

Redundancy Status: Down

Redundancy Mode: N/A

Pair Primary Status: Local Active

Pair Backup Status: Shutdown

Auto Revert: false

ADB Link To Mate Up: false

ADB Hello To Mate Up: false

Last Update: 04-Mar-2019 14:39:28

Fabric

Product	Fw-Version	Card Type	Slot	Serial #
NAB-0801ET-01-A	6.2.0.406	Network Acceleration Blade	1/1	S003000276
		in use by slot 1/1	1/2	
TRB-000000-02-A		Topic Routing Blade	1/3	P004045787
HBA-0204FC-02-A		Host Bus Adapter Blade	1/4	LFC0848B99409
ADB-000000-01-A		Assured Delivery Blade	1/5	S003000844
		empty	2/1	
		empty	2/2	
		empty	2/3	

- Host Name

The name of the host.
- Platform

The platform on which the broker is running.
- Chassis Product #

The product number of the chassis in which the broker is contained.
- Chassis Revision #

The revision number of the chassis.
- Chassis Serial #

The serial number of the chassis.
- Power Configuration

The power configuration used by the chassis.
- Operational Power Supplies

The number of available power supplies that are operational on the chassis.
- CPU 1

The name of the central processing unit (CPU 1) used by the broker.

CPU 2 The name of the central processing unit (CPU 2) used by the broker.

BIOS The basic input/output system used by the chassis.

Memory (KB)

Physical	Lists the Total amount, the Free amount, the Used amount, and the Used % of physical memory.
Swap	Lists the Total amount, the Free amount, the Used amount, and the Used % of swap memory.

Redundancy

These fields describe a fault tolerant pair of brokers.

Mate Router Name	If redundancy is configured, this field lists the redundant broker name (mate broker name).
Configuration Status	The status of the configuration for the backup broker.
Redundancy Status	The status of the redundant broker.
Redundancy Mode	Refer to Solace documentation for more information.
Primary Status	The status of the primary broker.
Backup Status	Refer to Solace documentation for more information.
Auto-Revert	Refer to Solace documentation for more information.
ADB Link Up	This check box is checked if a broker is set up to use guaranteed messaging and an Assured Delivery Blade (ADB) is set up and working correctly.
ADB Hello Up	Refer to Solace documentation for more information.

Fabric

Slot	Displays the slot number on the network switch.
Card Type	The type of card connected to the particular slot.
Product	The product associated with the particular slot.
Serial #	The serial number of the product.
Fw-Version	The firmware version of the product.

Broker Interface

This display lists all network interfaces on a selected broker, and shows network interface status, in/out throughput per second and additional detailed metrics.

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

You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.

Search by clicking the right side of a column heading/**Filter** to open the Search, Sort and Choose Columns dialog:

Filter

Show items with value that:

Sort Ascending

Sort Descending

Columns

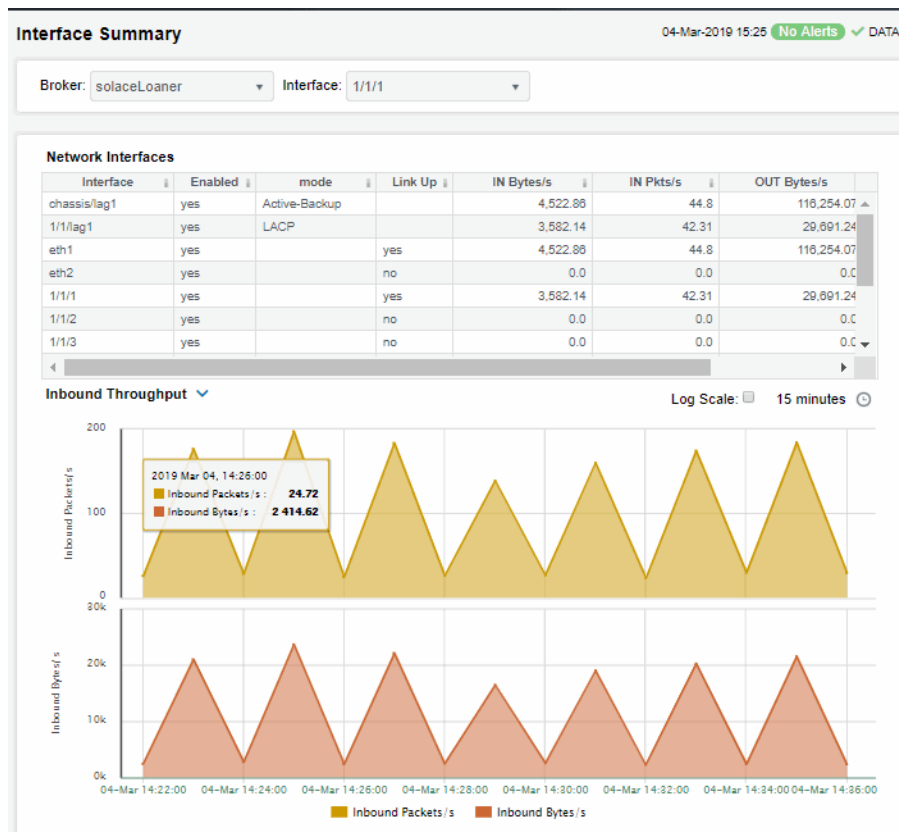
Contains

And

Contains

Filter

Clear



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

OUT Pkts/sec The number of outgoing packets per second.



Trend Graphs


Inbound Pkts/ sec Traces the number of incoming packets per second.

Outbound Bytes/sec Traces the number of bytes per second contained in the incoming messages.


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

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

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



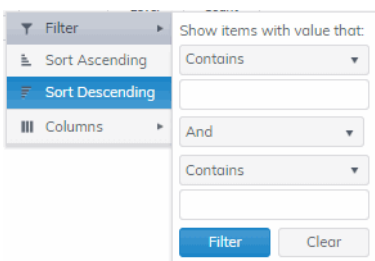
Toggle forward/backward in the trend graph per the period you choose (from the **Time range** drop-down menu) using arrows  .

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

Brokers Message Spool

Select a broker from the drop-down menu or select **All**. This display shows operational status and spooling performance metrics (if spooling is enabled on the broker) for one or all brokers.

Search by clicking the right side of a column heading/**Filter** to open the Search, Sort and Choose Columns dialog:



Refer to Solace documentation for details about data in this display.

Solace Broker Spool Table 04-Mar-2019 15:35 No Alerts DATA

Broker: - All -

Count: 1

Pool MB	Msg Spool Used By Queue	Msg Spool Used By DTE	Message Count % Usage	Delivered UnAcked Msgs % Usage	Ingress Flow Count	Outgress Flow Count
0.0	13	1	0.0	0.0	18	

Count	The number of brokers that are using spooling in the table.
Connection	The connection string associated with the broker.
Config Status	The message spool configuration status.
Operational Status	The operational status of the message spool.
Current Spool Usage (MB)	The current amount of spool used in megabytes on the broker (calculated by summing spool used for each endpoint).
Msg Spool Used By Queue	The amount of spool used by queue.
Msg Spool Used By DTE	The amount of spool used by DTE.
Message Count % Utilization	The percentage messages that use the message spool.
Delivered UnAcked Msgs % Utilization	The percentage of unacknowledged messages delivered from the message spool.
Ingress Flow Count	The current incoming flow count.
Ingress Flows Allowed	The number of incoming flows allowed.
Topic Subscriptions on Queue Used	The number of queue/topic subscriptions used.
Max Topic Subscriptions on Queue	The maximum number of queue/topic subscriptions available.
Sequenced Topics Used	The number of sequenced topics used.
Max Sequenced Topics	The maximum number of sequenced topics available.
Spool Files Used	The number of spool files used.
Spool Files Available	The maximum number of spool files available.

Spool Files % Utilization	The percentage of available spool files that have been used.
Active Disk Partition % Usage	The percentage of active disk partition that has been used.
Standby Disk Partition % Usage	The percentage of standby disk partition that has been used.
Disk Usage Current (MB)	The current amount of spool disk usage in megabytes.
Disk Usage Max (MB)	The maximum amount of spool disk usage in megabytes.
Transacted Sessions Used	The current number of transacted sessions.
Transacted Sessions Max	The maximum number of transacted sessions .
Transacted Session Count % Utilization	The percentage of transacted sessions that have been used.
Transacted Session Resource % Utilization	The percentage of transacted session resources that have been used.
Expired	When checked, performance data about the broker has not been received within the time specified.

CSPF Neighbors

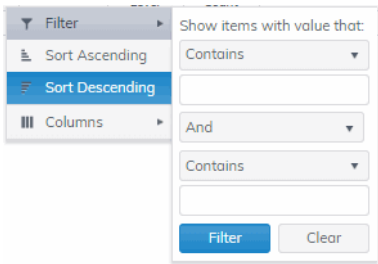
These displays provide detailed data and statuses for CSPF neighbor brokers. You can check trends on network traffic among CSPF neighbors. Note that these displays are empty if you are only monitoring Solace Cloud PubSub+ Brokers. Displays in this View are:

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

Neighbors Table

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

Search by clicking the right side of a column heading/**Filter** to open the Search, Sort and Choose Columns dialog:



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

Solace CSPF Neighbors Table

14-Aug-2018 16:16 No Alerts DATA

Broker:

More Columns

Show Ok Only:

Show: All

Neighbors: 4

Message Router	Name	Expired	State	Sent Msgs/s	Sent Bytes/s	Connections
VMR-112	ip-172-30-1-144		Ok	0.15	24.38	4
VMR-144	ip-172-30-1-112		Ok	0.2	0.0	4
VMR-144	ip-172-30-1-209		Ok	0.2	0.0	4
VMR-209	ip-172-30-1-144		Ok	0.17	29.41	4

- Neighbor Count:

The number of neighbor brokers connected to the selected Broker.
- Show:

OK

Select to *only* show neighbor brokers that are connected (**State** is **OK**). By default, this option is not selected (**all neighbor brokers are shown**).

Expired

Select to show *both* expired and non-expired neighbor brokers. By default, this option is not selected (only non-expired neighbor brokers are shown).

Table:
Each table row is a different neighbor broker.

- Broker

The name of the neighbor broker.
- State

The current state of the neighbor link.
- Up Time

The amount of time the broker has been up and running.
- Connections

The number of connections.
- Link Cost Actual

Refer to Solace documentation for more information.




Link Cost Configured	Refer to Solace documentation for more information.
Data Port	Refer to Solace documentation for more information.
Expired	When checked, performance data about the broker has not been received within the time specified.
Timestamp	The date and time the row of data was last updated.

CSPF Neighbors Diagram

Use this topology display to monitor the health of network components: Solace brokers, VMRs, servers and neighbor links. Quickly identify broker neighbors that are inactive and which resources their performance impacts.

Each node in the display is a Solace broker, VMR or server. If the value of the node's **Connected** column is **true**, then the node is green. If the value in its **Expired** column is **true**, it is dark gray. Otherwise the node is not connected and not expired, and the color is light gray.

No Alerts shows the number of current alerts associated with the objects and uses the following color code:

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

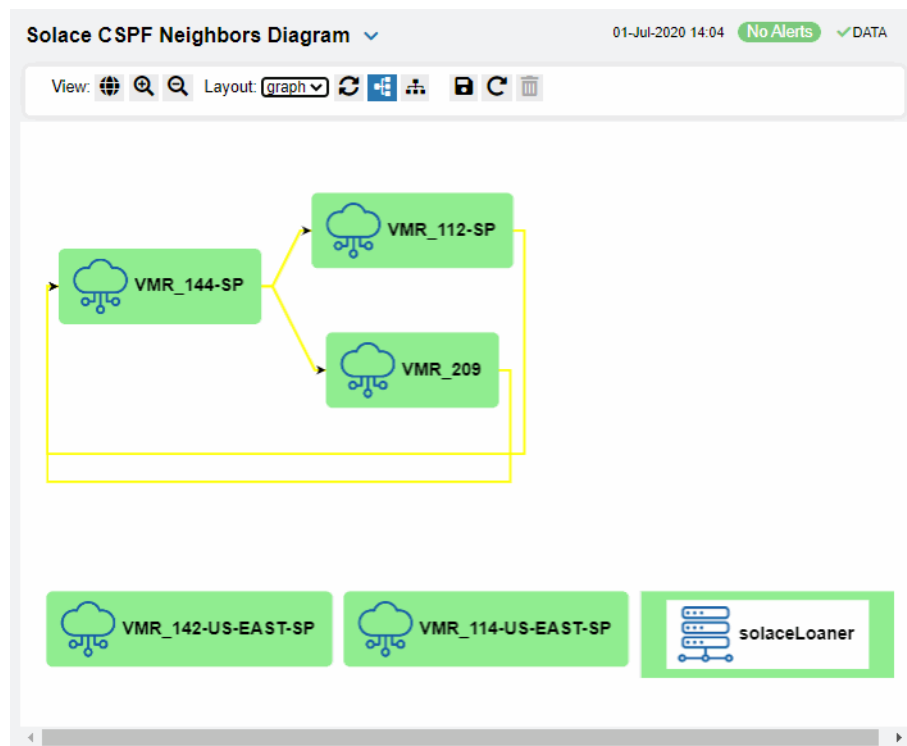
The lines connecting nodes represent neighbor links. If the value of the neighbor link **State**, is **OK** (which is shown in the ["Neighbors Table"](#)), then the neighbor link is green and if not the neighbor link yellow.

Drag and drop objects to arrange them on the screen (doing so does not logically impact Solace brokers, software or servers).

This display only shows items that are categorized as **Appliance**, **Appliance-Pair**, **VPN**, **ApplianceCapacity**, **Client**, **Bridge** or **Endpoint**. And when you click **No Alerts** to open the **Alerts Table by Component** display and investigate alerts, only alerts for items in these categories are listed.

You can mouse-over objects to see their **Host IP address** and **Platform**. Right-click on VMR objects and select **Open VMR UI** to open the Solace VMR login web page. Double-click to drill down and investigate in the ["Neighbors Table"](#).

(Please scroll down for details about using the tool bar.)



Use the tool bar to interact with the diagram:



- Toggle to include/exclude the overview (a thumbnail version of display elements that appears in the upper left margin of the display).



- Zoom In/Zoom Out



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



- Apply a vertical flow/horizontal flow (these options apply to tree and graph layout only).



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




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



- Delete the locally saved layout options.

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



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

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

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

To monitor network bridges and VPNs, see the [“Bridges Diagram”](#).

Neighbors Summary

View neighbor broker current configuration details and message throughput rates.

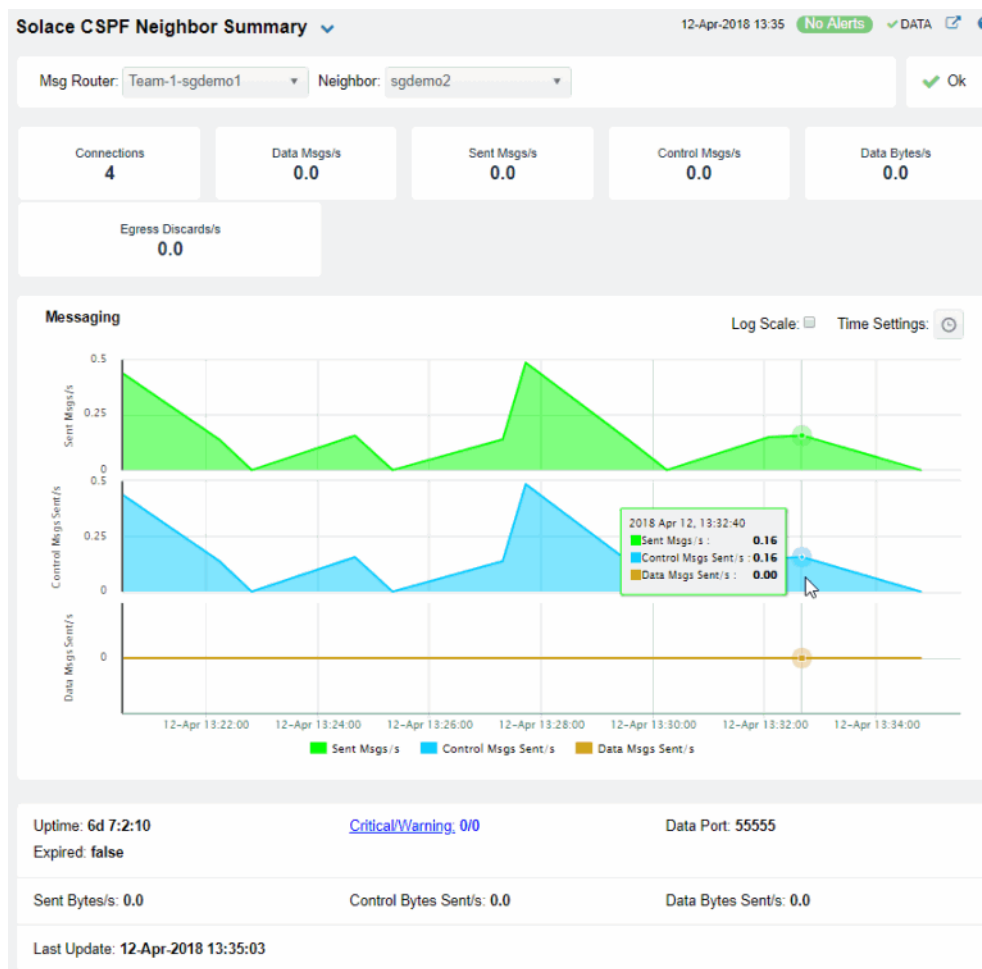
Select a broker and a neighbor broker from the drop down menus. Check message throughput rates to the neighbor broker, as well as neighbor **Up Time**, **State**, **Data Port**, number of connections and link costs.

You can hover over the metric cards to see more performance metrics and also drill down to see even more detail by clicking on them.

The bottom half of the display provides current and historical performance metrics for the selected broker. The trend graph traces the performance metric you select: **Message Flow** or **Throughput**.

You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.

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



Neighbor: Select the neighbor broker for which you want to show data in the display.

Connections The current number of connections.

Data Msgs/s Refer to Solace documentation for more information.

Sent Msgs/s Refer to Solace documentation for more information.

Control Msgs/s Refer to Solace documentation for more information.

Data Bytes/s Refer to Solace documentation for more information.



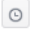
Egress Discards/s The total number of discarded messages sent from the selected **Broker** to the selected **Neighbor** broker since the broker was last started.


Trend Graphs



Traces the rates of messages sent from the selected **Broker** to the selected **Neighbor** broker.

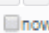
Sent Msgs/s Refer to Solace documentation for more information.

Control Msgs/s Refer to Solace documentation for more information.

- Discards/s** Traces the number of discarded messages sent, per second, from the selected **Broker** to the selected **Neighbor** broker.
- Log Scale** Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. **Log Scale** makes data on both scales visible by applying logarithmic values rather than actual values to the data.
- Time Settings** By default, the time range end point is the current time. To change the time range, click the **Time Settings**  and either:
- choose a **Time range** from 5 Minutes to 7 Days in the drop-down menu.
 - specify begin/end dates using the calendar .
 - specify begin/end time using the clock .



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

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

VPNs

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


- [“VPNs Heatmap” on page 99](#): A color-coded heatmap view of the current status of all VPNs configured on a specific broker.
- [“VPNs Table” on page 103](#): A tabular view of all available data for all VPNs configured on a specific broker.
- [“VPNs Summary” on page 106](#): Current and historical metrics for a single VPN.

VPNs Heatmap

View the status of all VPNs configured on a specific broker in a heatmap format, which allows you to quickly identify VPNs with critical alerts. Each rectangle in the heatmap represents a VPN. The rectangle color indicates the alert state and rectangle size represents the number of connections.

Select a broker from the **Broker** drop-down menu, or enter a search string in the **Filter VPN Name** field, and select a metric from the **Metric** drop-down menu. Use the **Show Operational Only** check-box ☒ to include or exclude non-operational VPNs in the heatmap. Use the **Log Scale** and **Auto Scale** check-boxes ☒ to apply log or auto scale. Use the **Show Broker** check-box ☒ to include or exclude broker names in the heatmap.

By default, this display shows **Alert Severity**, but you can mouse over a rectangle to see additional metrics. Drill down and investigate by clicking a rectangle in the heatmap to view details for the selected application in the “[VPNs Summary](#)” display.

**Operational**

When checked, only shows operational brokers.

Filter VPN Name

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

Metric

Choose a metric to view in the display.

Alert Severity


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

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

Alert Count


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

Connections

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


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

Subscriptions

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

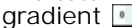
When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

Msgs Spooled

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


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

Total Msgs Rcvd

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

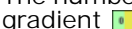
The **Auto** flag does not impact this metric.

Total Msgs Sent

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


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

Total Msgs/ sec Rcvd


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

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


Total Msgs/ sec Sent

The number of messages sent per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **SolVpnOutboundMsgRateHigh**. The middle value in the gradient bar indicates the middle value of the range. When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

Total Bytes/ sec Rcvd


The number of bytes contained in messages received per second. The color gradient  bar, populated by the current heatmap, shows the value/color mapping. The numerical values in the gradient bar range from 0 to the defined alert threshold of **SolVpnInboundByteRateHigh**. The middle value in the gradient bar indicates the middle value of the range. When **Auto** is checked, the numeric values in the color gradient bar show the range of the data being displayed rather than the default values. The middle value changes accordingly to indicate the color of the middle value of the range.

Total Bytes/ sec Sent

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


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

Direct Msgs/sec Rcvd

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


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

Direct Msgs/sec Sent




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

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

Total Inbound Discards

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

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

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

VPNs Table


View data shown in the “[VPNs Heatmap](#)” display, as well as additional details, in a tabular format. Use this display to view all available data for each VPN associated with a specific broker.

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

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

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

Search by clicking the right side of a column heading/**Filter** to open the Search, Sort and Choose Columns dialog:



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

Solace VPNs Table

04-Mar-2019 16:05 No Alerts DATA

Broker: - All - Less Columns

Show Operational Only: Show: All

Filter VPN Name: * VPNs: 25

Broker	VPN Name	Alert Level	Alert Count	Connections	Operational
solaceLoaner	BridgeBroker1	✓		3	✓
solaceLoaner	Broker1	✓		3	✓
solaceLoaner	Broker10	✓		3	✓
solaceLoaner	Broker2	✓		3	✓
solaceLoaner	Broker3	✓		3	✓
solaceLoaner	Broker4	✓		3	✓
solaceLoaner	Broker5	✓		3	✓
solaceLoaner	Broker6	✓		3	✓
solaceLoaner	Broker7	✓		3	✓
solaceLoaner	Broker8	✓		3	✓
solaceLoaner	Broker9	✓		3	✓
solaceLoaner	default	✓		0	
solaceLoaner	stats	✓		0	
solaceLoaner	statspump	✓		1	✓
solaceLoaner	white-test	✓		1	✓

- Broker

VPN Name

Alert Level

Alert Count

Connections
- The name of the broker.

The name of the VPN.

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.

The total number of active alerts for the VPN.

The total number of connections for the VPN.
- 104
- Solace PubSub+ Monitor User's Guide

Operational	When checked, this status indicates that the VPN is enabled and is operating normally.
Total Unique Subscriptions	The total number of unique subscriptions to the VPN.
Total Client Messages Rcvd	The total number of messages received from clients connected to the VPN.
Total Client Messages Sent	The total number of messages sent to clients connected to the VPN.
Total Client Bytes Rcvd	The total number of bytes contained in messages received from clients connected to the VPN.
Total Client Bytes Sent	The total number of bytes contained in messages sent to clients connected to the VPN.
Total Client Msgs/sec Rcvd	The total number of messages received per second from clients connected to the VPN.
Total Client Msgs /sec Sent	The total number of messages sent per second to clients connected to the VPN.
Total Client Bytes/sec Rcvd	The total number of bytes contained in messages received per second from clients connected to the VPN.
Total Client Bytes/sec Sent	The total number of bytes contained in messages sent per second to clients connected to the VPN.
Client Direct Msgs Rcvd	The total number of direct messages received from clients connected to the VPN.
Client Direct Msgs Sent	The total number of direct messages sent to clients connected to the VPN.
Client Direct Bytes Rcvd	The total number of bytes contained in direct messages received from clients connected to the VPN.
Client Direct Bytes Sent	The total number of bytes contained in direct messages sent to clients connected to the VPN.
Client Direct Msgs/sec Rcvd	The total number of direct messages received per second from clients connected to the VPN.
Client Direct Msgs/sec Sent	The total number of direct messages sent per second to clients connected to the VPN.
Client Direct Bytes/sec Rcvd	The total number of bytes contained in the direct messages received per second from clients connected to the VPN.
Client Direct Bytes/sec Sent	The total number of bytes contained in the direct messages sent per second to clients connected to the VPN.
Client NonPersistent Msgs Rcvd	The total number of non-persistent messages received from clients connected to the VPN.
Client NonPersistent Msgs Sent	The total number of non-persistent messages sent to clients connected to the VPN.
Client NonPersistent Bytes Rcvd	The total number of bytes contained in the non-persistent messages received from clients connected to the VPN.
Client NonPersistent Bytes Sent	The total number of bytes contained in the non-persistent messages sent per second to clients connected to the VPN.
Client NonPersistent Msgs/sec Rcvd	The total number of non-persistent messages received per second from clients connected to the VPN.
Client NonPersistent Msgs/sec Sent	The total number of non-persistent messages sent per second to clients connected to the VPN.

Client NonPersistent Bytes/sec Rcvd	The total number of bytes contained in the non-persistent messages received per second from clients connected to the VPN.
Client NonPersistent Bytes/sec Sent	The total number of bytes contained in the non-persistent messages sent per second to clients connected to the VPN.
Client Persistent Msgs Rcvd	The total number of persistent messages received from clients connected to the VPN.
Client Persistent Msgs Sent	The total number of persistent messages sent to clients connected to the VPN.
Client Persistent Bytes Rcvd	The total number of bytes contained in persistent messages received from clients connected to the VPN.
Client Persistent Bytes Sent	The total number of bytes contained in persistent messages sent to clients connected to the VPN.
Client Persistent Msgs/sec Rcvd	The total number of persistent messages received per second from clients connected to the VPN.
Client Persistent Msgs/sec Sent	The total number of persistent messages sent per second to clients connected to the VPN.
Client Persistent Bytes/sec Rcvd	The total number of bytes contained in the persistent messages received per second from clients connected to the VPN.
Client Persistent Bytes/sec Sent	The total number of bytes contained in the persistent messages sent per second to clients connected to the VPN.
Total In Discards	The total number of discarded incoming messages.
Total In Discards/sec	The number of discarded incoming messages per second.
Total Out Discards	The total number of discarded outgoing messages.
Total Out Discards/sec	The number of discarded outgoing messages per second.
Max Spool Usage (MB)	The maximum amount of disk storage (in megabytes) that can be consumed by all spooled message on the VPN.
Authentication Type	The defined authentication type on the VPN.
Expired	When checked, performance data about the broker has not been received within the time specified.
Time Stamp	The date and time the row data was last updated.

VPNs Summary

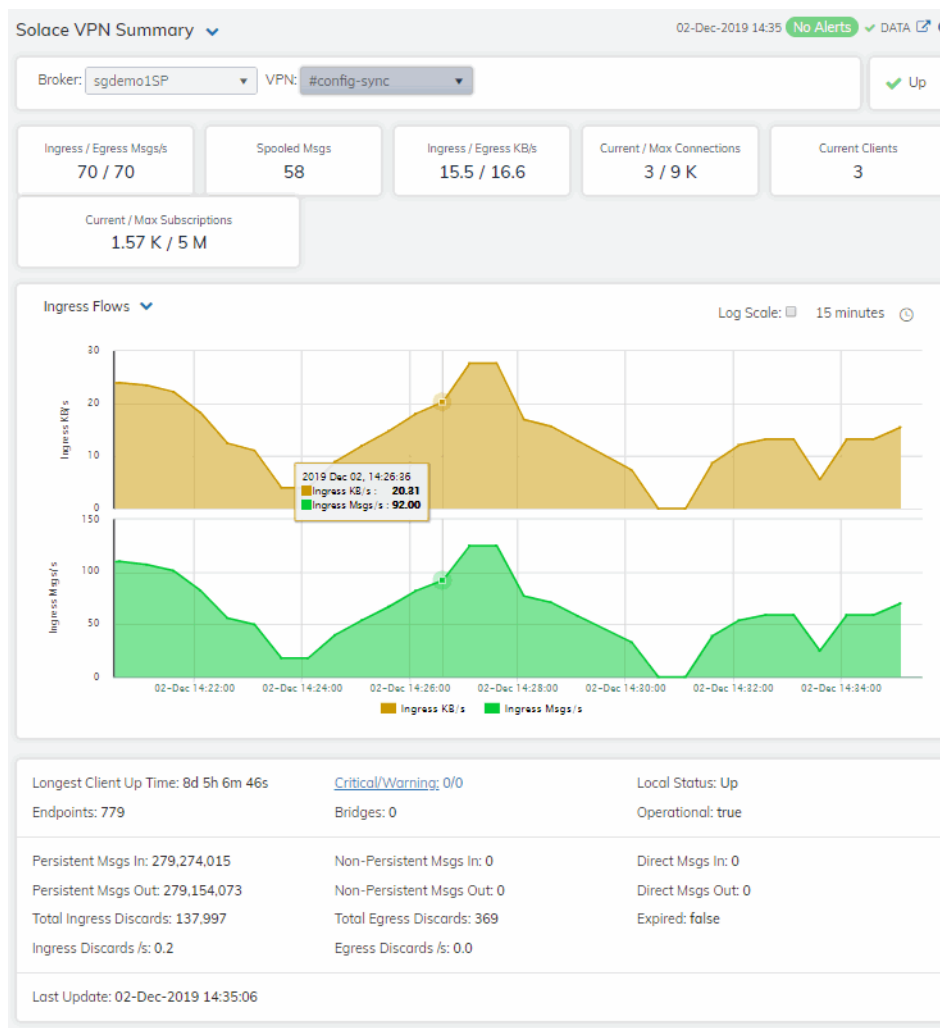
View neighbor broker current configuration details and message throughput rates.

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

You can hover over the metric cards to see more performance metrics and also drill down to see even more detail by clicking on them.

The bottom half of the display provides current and historical performance metrics for the selected broker. The trend graph traces the performance metric you select: **Ingress Flows**, **Egress Flows** or **Spool Msgs**.

You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.



Alerts

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

Up

Inbound/Outbound Msgs/s	The number of inbound/outbound messages per second.
Spooled Msgs	The number of spooled messages.
Inbound/Outbound KB/s	The number of inbound/outbound messages in KBs per second.
Current/Max Connections	The total number of current connections / maximum number of supported connections for the VPN.
Current Clients	The number of connected clients.

Current/Max Subscriptions The total number of current subscribers and maximum number of supported subscribers for the VPN.

Inbound Msgs/s Trend Graphs

Traces the sum of inbound message processing for the selected VPN.

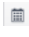

- **Spooled Msgs:** The number of spooled messages for the VPN.
- **Client Msgs/sec:** The rate of incoming messages (per second) from client.
- **Direct Client Msgs/sec:** The rate of direct incoming messages (per second) from the direct client.


Log Scale



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


Time Settings

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

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



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

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

Longest Client Up Time

The number of days, hours and minutes for the longest, currently active, client connection.

Endpoints

The number of endpoints.

Persistent Msgs In

The total number of incoming persistent messages.

Persistent In Msgs/s

The number of incoming persistent messages per second.

Persistent Msgs Out

The total number of outgoing persistent messages.

Persistent Out Msgs/s

The number of outgoing persistent messages per second.

Total In Discards

The total number of incoming messages that were discarded.

Total In Discards/sec

The total number of incoming messages that were discarded, per second.

Critical/Warning

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

Bridges

The number of bridges.

Non-Persistent Msgs In

The total number of incoming non-persistent messages.

Non-Persistent In Msgs/s	The number of incoming non-persistent messages per second.
Non-Persistent Msgs Out	The total number of outgoing non-persistent messages.
Non-Persistent Out Msgs/s	The number of outgoing non-existent messages per second.
Total Out Discards	The total number of outgoing messages that were discarded.
Total Out Discards/sec	The total number of outgoing messages that were discarded, per second.
Direct Msgs In	The total number of incoming direct messages.
Direct In Msgs/s	The number of incoming direct messages per second.
Direct Msgs Out	The total number of outgoing direct messages.
Direct Out Msgs/s	The number of outgoing direct messages per second.
Expired	When true , performance data about the VPN has not been received within the time specified.
Last Update	The date and time of the last data update.

Clients

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

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

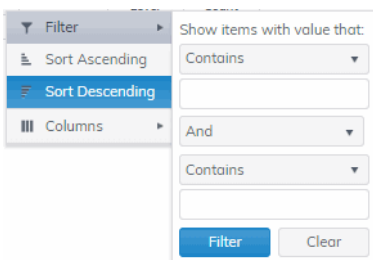
Clients Table

View VPN clients configured on all brokers, a single broker, all VPNs or a single VPN. Each table row is a different VPN client connection. Use the drop-down menus to show **All**, **Expired** or **Unexpired** clients as well as **All**, **Internal** or **Primary** clients (processes that run on the broker under the Solace OS). Enter a string for **Filter Client Name** to show only clients with this string in their name.

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

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

Search by clicking the right side of a column heading/**Filter** to open the Search, Sort and Choose Columns dialog:



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

Solace Broker Clients Table

05-Mar-2019 08:38 No Alerts DATA

Broker: - All - VPN: - All - Less Columns

Show Type: All Show: All Filter Client Name: *

Clients: 101

Broker	VPN	Client Name	Alert Level	Alert Count	Slow Subscriber
solaceLoaner	BridgeBroker1	#bridge/local/B114toSolDemo/solace/8798/16	✓		
solaceLoaner	BridgeBroker1	#bridge/local/B142toSolDemo/solace/8798/15	✓		
solaceLoaner	BridgeBroker1	#client	✓		
solaceLoaner	Broker1	#bridge/local/testBridgeToNoWhere/solace/8798/9	✓		
solaceLoaner	Broker1	#bridge/remote/B1_to_B2/v/solace/8796/0	✓		
solaceLoaner	Broker1	#client	✓		
solaceLoaner	Broker1	S-HOST10/5236/#00010001	✓		
solaceLoaner	Broker10	#bridge/local/B112toSolDemo/solace/8798/14	✓		
solaceLoaner	Broker10	#bridge/local/B144toSolDemo/solace/8798/12	✓		
solaceLoaner	Broker10	#bridge/local/B209toSolDemo/solace/8798/13	✓		
solaceLoaner	Broker10	#client	✓		
solaceLoaner	Broker10	S-HOST10/5152/#00010001	✓		
solaceLoaner	Broker10	S-HOST10/5448/#00010001	✓		
solaceLoaner	Broker2	#bridge/local/B1_to_B2/solace/8798/10	✓		
solaceLoaner	Broker2	#client	✓		
solaceLoaner	Broker2	S-HOST10/5212/#00010001	✓		
solaceLoaner	Broker3	#bridge/local/Bridge_loanerToVMR144/solace/879	✓		

Page 1 of 3 1 - 40 of 101 items

Broker	Lists the name of the selected broker.
VPN	Lists the name of the selected VPN.
Client Name	The name of the client.
Alert Level	The maximum level of alerts in the row: <div><div>●</div> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</div> <div><div>●</div> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</div> <div><div>●</div> Green indicates that no metrics have exceeded their alert thresholds.</div>
Alert Count	Total number of alerts for the client.
Slow Subscriber	This check box will be checked if the client consistently fails to consume their messages at the offered rate (which causes their egress queues to fill up).
Total Egress Flows	The total number of outgoing flows.
Total Ingress Flows	The total number of incoming flows.
Subscriptions	The total number of subscriptions.
Subscription Msgs Rcvd	The total number of messages received from subscriptions.

Subscription Msgs Sent	The total number of messages sent from subscriptions.
Type	Lists the type of alert.
Uptime	Lists the amount of time the client has been up and running.
Client ID	Lists the client ID.
Client UserName	Lists the user name for the client.
Client Address	The IP Address of the client.
Client Profile	The client profile that is assigned to the client.
ACL Profile	The access control list profile to which the client is assigned.
Description	Lists a description of the client.
Platform	Lists the platform of the client.
Software Version	The version of the platform.
Total Flows Out	The total number of outbound message flows for the client.
Total Flows In	The total number of inbound message flows for the client.
# Subscriptions	The number of subscribers connected to the client.
Add Sub Msgs Rcvd	The number of Add Subscription messages received.
Add Sub Msgs Sent	The number of Add Subscription Messages sent.
Already Exists Msgs Sent	Refer to Solace documentation for more information.
Assured Ctrl Msgs Rcvd	Refer to Solace documentation for more information.
Assured Ctrl Msgs Sent	Refer to Solace documentation for more information.
Total Client Msgs Rcvd	The total number of messages received by the client.
Total Client Msgs Sent	The total number of messages sent by the client.
Total Client Bytes Rcvd	The total number of bytes contained within the messages received by the client.
Total Client Bytes Sent	The total number of bytes contained within the messages sent by the client.
Total Client Msgs Rcvd/sec	The total number of messages received per second by the client.
Total Client Msgs Sent/sec	The total number of messages sent per second by the client.
Total Client Bytes Rcvd/sec	The total number of bytes contained within the messages received per second by the client.
Total Client Bytes Sent/sec	The total number of bytes contained within the messages sent per second by the client.
Ctl Bytes Rcvd	The number of control data bytes received by the client.
CTL Bytes Sent	The number of control data bytes sent by the client.

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

Client Persistent Msgs Sent/sec	The number of persistent messages sent per second by the client.
Client Persistent Bytes Rcvd/sec	The number of bytes contained within the persistent messages received per second by the client.
Client Persistent Bytes Sent/sec	The number of bytes contained within the persistent messages sent per second by the client.
Denied Dup Clients	Refer to Solace documentation for more information.
Denied Subscribe Permission	The number of denied subscription requests due to improper permissions.
Denied Subscribe Topic-ACL	The number of denied subscriptions to topics due to the fact that the client requesting was not on the Access Control List.
Denied Unsubscribe Permission	The number of denied unsubscribe requests due to improper permissions.
Denied Unsubscribe Topic-ACL	The number of denied unsubscribe requests to topics due to the fact that the client requesting was not on the Access Control List.
DTO Msgs Rcvd	The number of Deliver-To-One messages received by the client.
Egress Compressed Bytes	The number of compressed bytes contained within outgoing messages.
Ingress Compressed Bytes	The number of compressed bytes contained within incoming messages.
Total Ingress Discards	The total number of discarded incoming messages.
Total Egress Discards	The total number of discarded outgoing messages.
Total Ingress Discards/sec	The total number of discarded incoming messages per second.
Total Egress Discards/sec	The total number of discarded outgoing messages per second.
Keepalive Msgs Rcvd	The number of Keepalive messages received by the client.
Keepalive Msgs Sent	The number of Keepalive messages sent by the client.
Large Msgs Rcvd	The number of large messages received by the client.
Login Msgs Rcvd	The number of login message received by the client.
Max Exceeded Msgs Sent	The number of responses sent by the client informing the connected broker(s) that the number of the message(s) sent exceeded the maximum allowed.
Not Enough Space Msgs Sent	The number of responses sent by the client informing the connected broker(s) that the size of the message(s) sent exceeded the maximum allowable size, or that the message caused the client's Local Spool Quota to exceed the maximum amount of space.
Not Found Msgs Sent	Refer to Solace documentation for more information.
Parse Error on Add Msgs Sent	Refer to Solace documentation for more information.
Parse Error on Remove Msgs Sent	Refer to Solace documentation for more information.
Remove Subscription Msgs Rcvd	The number of remove subscription requests received by the client.

Remove Subscription Msgs Sent	The number of remove subscription requests sent by the client.
Subscribe Client Not Found	The number of subscription requests for clients that were not found.
Unsubscribe Client Not Found	The number of unsubscribe requests for clients that were not found.
Update Msgs Rcvd	Refer to Solace documentation for more information.
Update Msgs Sent	Refer to Solace documentation for more information.
Expired	When checked, performance data about the client has not been received within the time specified.
Timestamp	The date and time the row of data was last updated.

Client Summary

View current and historical performance and utilization metrics for a single VPN client.

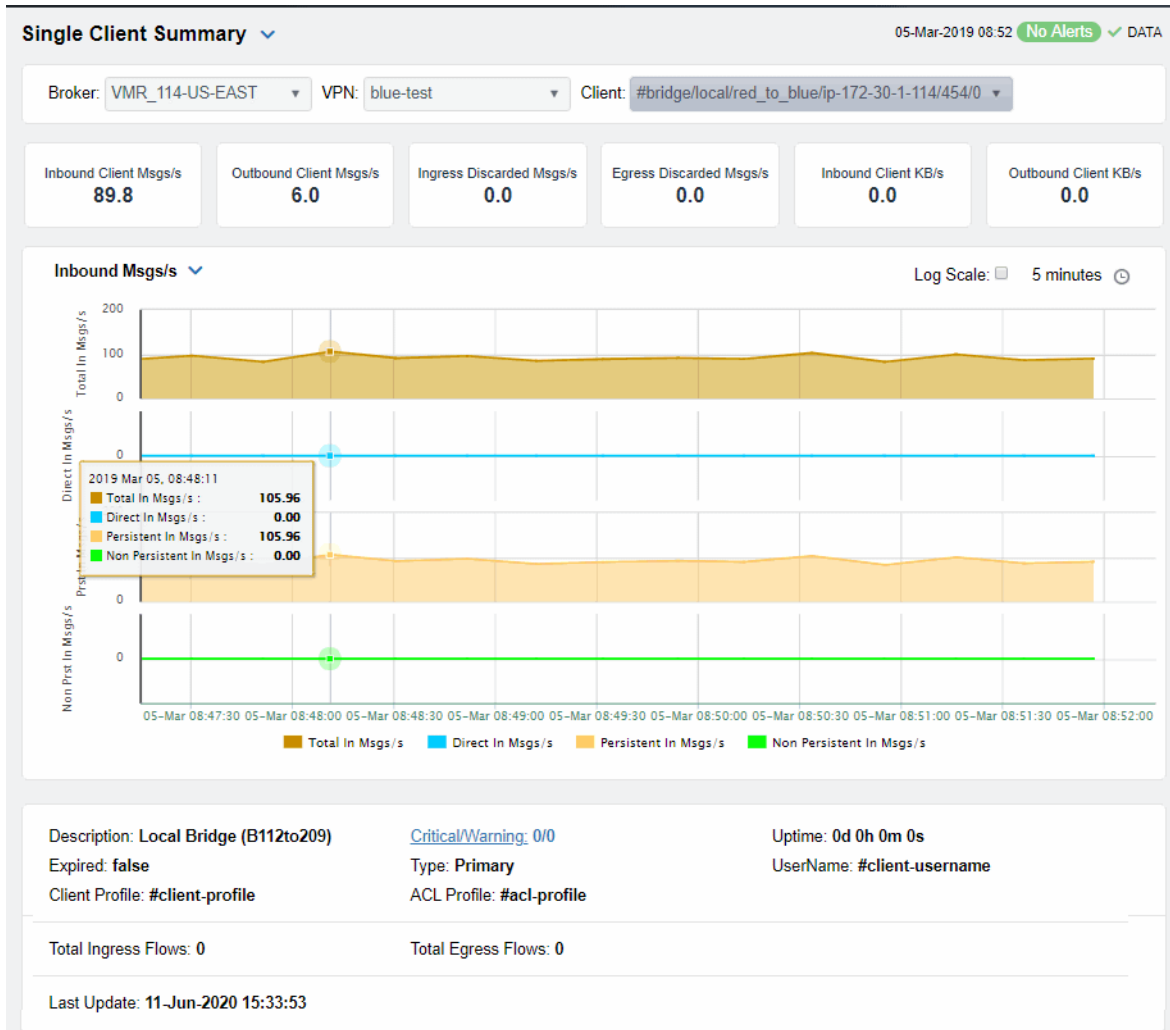
Select a broker, VPN and client from the drop-down menus. You can view the **Client Type**, the **User Name**, the **Client ID**, the associated **Platform**, the current **Up Time**, and additional information specific to the client. You can also view the total number of incoming and outgoing messages,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

You can hover over the metric cards to see more performance metrics and also drill down to see even more detail by clicking on them.

The bottom half of the display provides current and historical performance metrics for the selected broker. The trend graph traces the performance metric you select: **Ingress Flows** or **Egress Flows**.

You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.

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



Inbound Client Msgs /sec

The number of incoming client messages per second.

Outbound Client Msgs /sec

The number of outgoing client messages per second.

Ingress Discarded Msgs /sec

The number of discarded ingress messages per second.

Egress Discarded Msgs /sec

The number of discarded egress messages per second.

Inbound Client KB/sec

The amount of incoming data from the client in KBs per second.

Outbound Client KB/sec

The amount of outgoing data for the client in KBs per second.

Trend Graphs

Traces the sum of message processing for the selected client.



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


Log Scale



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

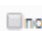
Time Settings

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

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



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

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

Description

The description of the client.

Expired

When checked, performance data about the broker has not been received within the time specified.

Client Profile

The client profile that is assigned to the client.

Total Ingress Flows

The number of inflows coming to the client.

Persistent Msgs In/sec

The number of persistent incoming messages per second.

Persistent Msgs Out/sec

The number of persistent outgoing messages per second.

Last Update

The date and time of the last data update.

Critical/Warning

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

Non Persistent Msgs In/sec

The number of non-persistent incoming messages per second.

NonPersistent Msgs Out/sec

The number of non-persistent outgoing messages per second.

Uptime

If the VPN's **Local Status** is **Up**, this field displays the length of time that the VPN has been up and running.

Username	The client's user name.
Direct In Msgs /sec	The number of non-persistent incoming messages per second.
Direct Out Msgs /sec	The number of non-persistent outgoing messages per second.

Bridges

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

- **"Bridges Table"**: A tabular view of all available process performance data for all bridges configured on a VPN.
- **"Bridges Diagram"**: Topological view of Solace network bridges that shows bridge broker connections and health status and allows you to open the Solace PubSub+ Manager.
- **"Bridge Summary"**: Current and historical metrics for a single bridge.

Bridges Table

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

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

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

Search by clicking the right side of a column heading/**Filter** to open the Search, Sort and Choose Columns dialog:

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.

Broker	Local VPN	Bridge Name	Remote VPN	Remote Router
VMR_114-US-EAST	blue-test	#bridge/v:solace/BridgeBroker1/15	BridgeBroker1	v:solace
VMR_114-US-EAST	blue-test	red_to_blue	red-test	v:ip-172-30-1-1

Broker

Displays the name of the broker

Local VPN

The name of the local VPN.

Bridge Name

The name of the bridge.

Alert Level

The current level of alerts in the row.

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

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

● Green indicates that no metrics have exceeded their alert thresholds.

Alert Count

The total number of active alerts for the process.

Remote VPN

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

Remote Router

The name of the remote broker.

Admin State

Indicates whether the bridge has been administratively enabled (via SolAdmin or the command line interface).

Inbound Operational State

The current inbound operational status of the bridge. (The administrator can turn off a bridge's input or output for maintenance or other reasons.)

Outbound Operational State

The current outbound operational status of the bridge. (The administrator can turn off a bridge's input or output for maintenance or other reasons.)

Queue Operational State

The current operational status of the queue.

Connection Establisher	Indicates whether the administrator created and configured the bridge directly on the broker using SolAdmin or the command line interface, or indirectly from another broker.
Redundancy	Displays whether the bridge is the primary bridge, the backup bridge, the static bridge (default bridge used when no other bridge is available), or whether it is the only bridge available (none).
Uptime	The current amount of time in which the bridge has been up and running.
Client Name	The name of the client.
Connected Via Addr	The local IP address and port used for the bridge.
Connected Via Interface	The name of the network interface used for the bridge.
Client Direct Bytes Rcvd	The number of bytes contained within direct messages received by the client via the bridge.
Client Direct Bytes/sec Rcvd	The number of bytes contained within direct messages received per second by the client via the bridge.
Client Direct Bytes Sent	The number of bytes contained within direct messages sent by the client via the bridge.
Client Direct Bytes/sec Sent	The number of bytes contained within direct messages sent per second by the client via the bridge.
Client Direct Msgs/sec Rcvd	The number of bytes contained within direct messages received per second by the client via the bridge.
Client Direct Msgs Sent	The number of direct messages sent by the client via the bridge.
Client Direct Msgs/sec Sent	The number of direct messages sent per second by the client via the bridge.
Client NonPersistent Bytes Rcvd	The number of bytes contained within non-persistent messages received by the client via the bridge.
Client NonPersistent Bytes/sec Rcvd	The number of bytes contained within non-persistent messages received per second by the client via the bridge.
Client NonPersistent Bytes Sent	The number of bytes contained within non-persistent messages sent by the client via the bridge.
Client NonPersistent Bytes/sec Sent	The number of bytes contained within non-persistent messages sent per second by the client via the bridge.
Client NonPersistent Msgs Rcvd	The number of non-persistent messages received by the client via the bridge.
Client NonPersistent Msgs/sec Rcvd	The number of non-persistent messages received per second by the client via the bridge.
Client NonPersistent Msgs Sent	The number of non-persistent messages sent by the client via the bridge.
Client NonPersistent Msgs/sec Sent	The number of non-persistent messages sent per second by the client via the bridge.
Client Persistent Bytes Rcvd	The number of bytes contained within persistent messages received by the client via the bridge.
Client Persistent Bytes/sec Rcvd	The number of bytes contained within persistent messages received per second by the client via the bridge.
Client Persistent Bytes Sent	The number of bytes contained within persistent messages sent by the client via the bridge.

Client Persistent Bytes/sec Sent	The number of bytes contained within persistent messages sent per second by the client via the bridge.
Client Persistent Msgs Rcvd	The number of persistent messages received by the client via the bridge.
Client Persistent Msgs /sec Rcvd	The number of persistent messages received per second by the client via the bridge.
Client Persistent Msgs Sent	The number of persistent messages sent by the client via the bridge.
Client Persistent Msgs/sec Sent	The number of persistent messages sent per second by the client via the bridge.
Total Client Bytes Rcvd	The number of bytes contained within all messages received by the client via the bridge.
Total Client Bytes/sec Rcvd	The number of bytes contained within all messages received per second by the client via the bridge.
Total Client Bytes Sent	The number of bytes contained within all messages sent by the client via the bridge.
Total Client Bytes/sec Sent	The number of bytes contained within all messages sent per second by the client via the bridge.
Total Client Msgs Rcvd	The total number of all messages received by the client via the bridge.
Total Client Msgs/sec Rcvd	The total number of all messages received per second by the client via the bridge.
Total Client Msgs Sent	The total number of all messages sent by the client via the bridge.
Total Client Msgs/sec Sent	The total number of all messages sent per second by the client via the bridge.
Total Out Discards	The total number of discarded outgoing messages sent by the client via the bridge.
Total Out Discards/sec	The total number of discarded outgoing messages sent per second by the client via the bridge.
Total In Discards	The total number of discarded incoming messages received by the client via the bridge.
Total In Discards/sec	The total number of discarded incoming messages received per second by the client via the bridge.
Expired	When checked, performance data about the broker has not been received within the time specified.
Timestamp	The date and time the row of data was last updated.

Bridges Diagram

Use this topology view to monitor the health of your network bridges and VPNs. Quickly identify bridge and VPN connections, their health status and which resources their performance impacts. Open the Solace PubSub+ Manager by right-clicking on a router and selecting **Launch PubSub+ Manager**.

Drag and drop objects to arrange them on the screen (doing so does not logically impact the network bridges and VPNs). Arrows show the connections between VPNs and bridges.

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

- Red indicates that the object has one or more alerts in a critical state.
- Yellow indicates that the object has one or more alerts in a warning state.

● Green indicates that there are no alerts on the object.

● Gray indicates that the object is off-line.

Save: Saves the arrangement of the objects.


Restore: Returns objects to their previous positions.

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

Look at the miniature view in (upper left) to see all objects in either layout. Or zoom into the display using **Ctrl**+/- or **Ctrl**+ mouse wheel.

Drill down to investigate in the [“Bridges Table”](#).

To monitor network brokers, VMRs and servers, see the [“CSPF Neighbors Diagram”](#).



Bridge Summary

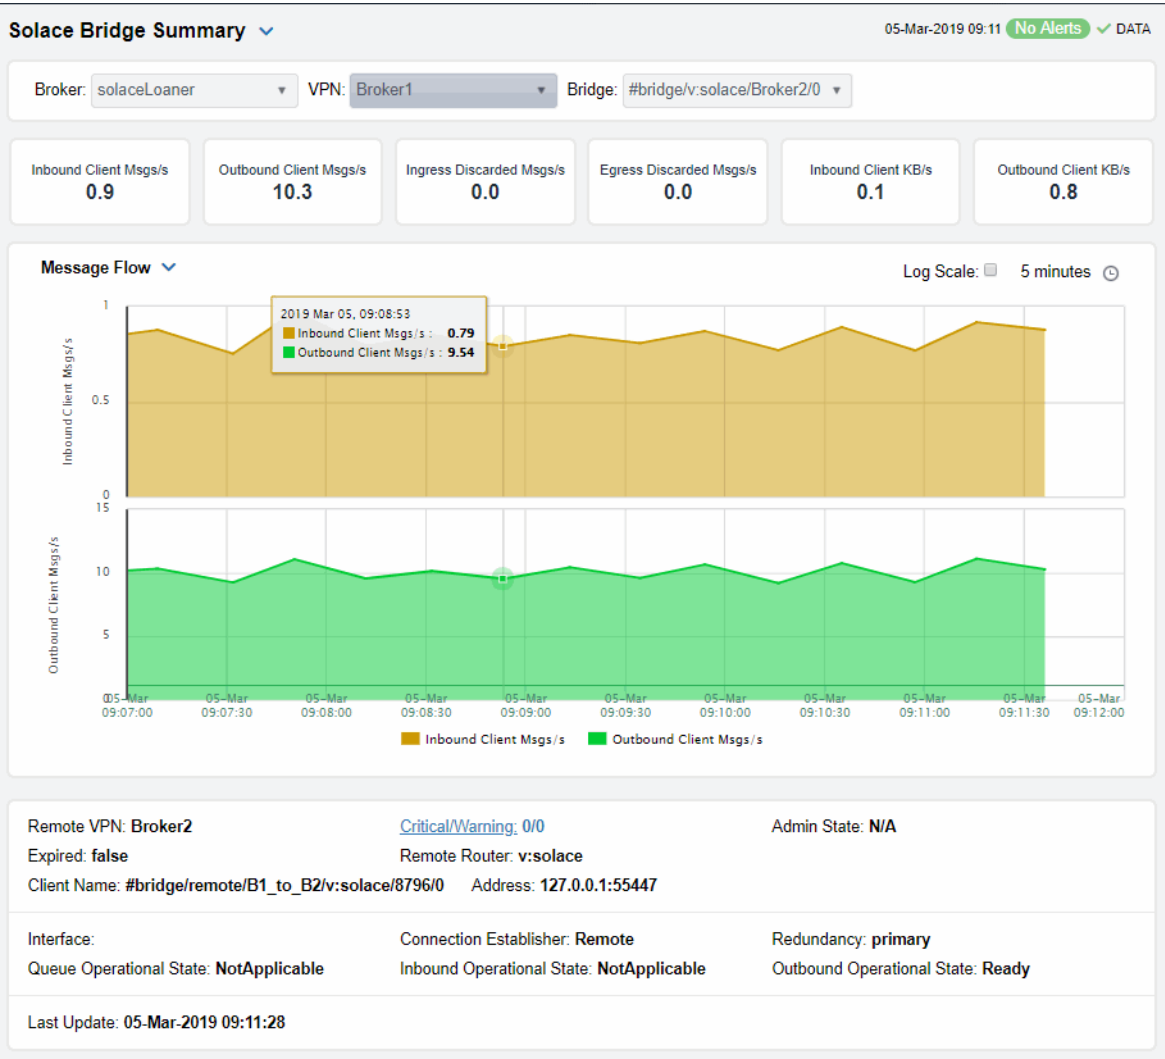
View current and historical performance and utilization metrics for a particular bridge on a VPN.

Select a broker, a VPN, and a bridge from the drop-down menus. Metric cards at the top of the displays show **Inbound and Outbound Client Messages per second**, **Ingress and Egress Discarded Messages**, and **Ingress and Egress KBs per second**.

You can hover over the metric cards to see more performance metrics and also drill down to see even more detail by clicking on them.

The trend graph traces current and historical performance metrics for the selected broker. The trend graph traces the performance metric you select: **Message Flow** or **Throughput**.

You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.



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

Messages Flow Trend Graphs

Traces the sum for the selected client.



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


Log Scale



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

Time Settings

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

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



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

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

Remote VPN

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

Expired

When true, performance data about the bridge has not been received within the time specified.

Address

The IP address.

Interface

The interface ID.

Queue Operational State

Refer to Solace documentation for more information.

Last Update

The date and time of the last data update.

Critical/Warning

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

Remote Router

The remote broker.

Conn Establisher

Refer to Solace documentation for more information.

Inbound Operational State

The current inbound operational status of the bridge. (The administrator can turn off a bridge's input or output for maintenance or other reasons.)

Admin State

Indicates whether the bridge has been administratively enabled (via SolAdmin or the command line interface).

Client Name

The name of the client.

Redundancy

Indicates whether the bridge is the **primary** bridge, the **backup** bridge, the **static** bridge (default bridge used when no other bridge is available), or whether it is the only bridge available (**none**).

Outbound Op State

The current outbound operational status of the bridge. (The administrator can turn off a bridge's input or output for maintenance or other reasons.)

Endpoints

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

- ["Endpoints Table"](#)
- ["Endpoint Summary"](#)

Endpoints Table

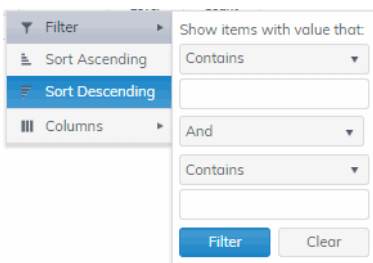
View all endpoints configured on a VPN. Each row in the table lists the details for a specific endpoint.

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

Select a broker and VPN from the drop-down menus. Filter the table using the **Show Ingress Config Status Down Only** check-box ☒ , use the **Show** drop-down menu to include **All**, **Expired** or **Unexpired** and use the **Type** drop-down menu to see **All**, **Queues Only** or **Topic Endpoint Only**.

When queue and topic from the same broker and VPN share the same name, **-TE** is appended to the topic name.

Search by clicking the right side of a column heading/**Filter** to open the **Search, Sort and Choose Columns** dialog:



You can click a column header to sort column data in numerical or alphabetical order, or double-click a row to drill down and investigate in the [“Endpoint Summary”](#) display.

Solace Endpoints Table 12-Jun-2020 15:51 No Alerts DATA

Broker: - All - VPN: - All - More Columns

Show Ingress Config Status Down Only: ☐ Show: All Type: All

Filter Endpoint Name : Endpoints: 38

Broker	VPN	Endpoint Name	Alert Level	Q/T Type	Endpoint Type	Durable
solaceLoaner	BridgeBroker1	BidgeBroker1Queue	✓	Q	Primary	✓
solaceLoaner	Broker1	bridgeq	✓	Q	Primary	✓
solaceLoaner	Broker1	q1	✓	Q	Primary	✓
solaceLoaner	Broker10	Broker10Queue	✓	Q	Primary	✓
solaceLoaner	Broker10	Broker10Topic1	✓	T	Primary	✓
solaceLoaner	Broker2	broker2q1	✓	Q	Primary	✓
solaceLoaner	Broker2	msg_from_b1	✓	Q	Primary	✓
solaceLoaner	Broker3	Broker3Queue	✓	Q	Primary	✓
solaceLoaner	Broker4	Broker4Queue	✓	Q	Primary	✓
solaceLoaner	Broker5	Broker5Queue	✓	Q	Primary	✓
solaceLoaner	Broker6	Broker6Queue	✓	Q	Primary	✓
solaceLoaner	Broker7	Broker7Queue	✓	Q	Primary	✓
solaceLoaner	Broker8	Broker8Queue	✓	Q	Primary	✓
solaceLoaner	Broker9	Broker9Queue	✓	Q	Primary	✓

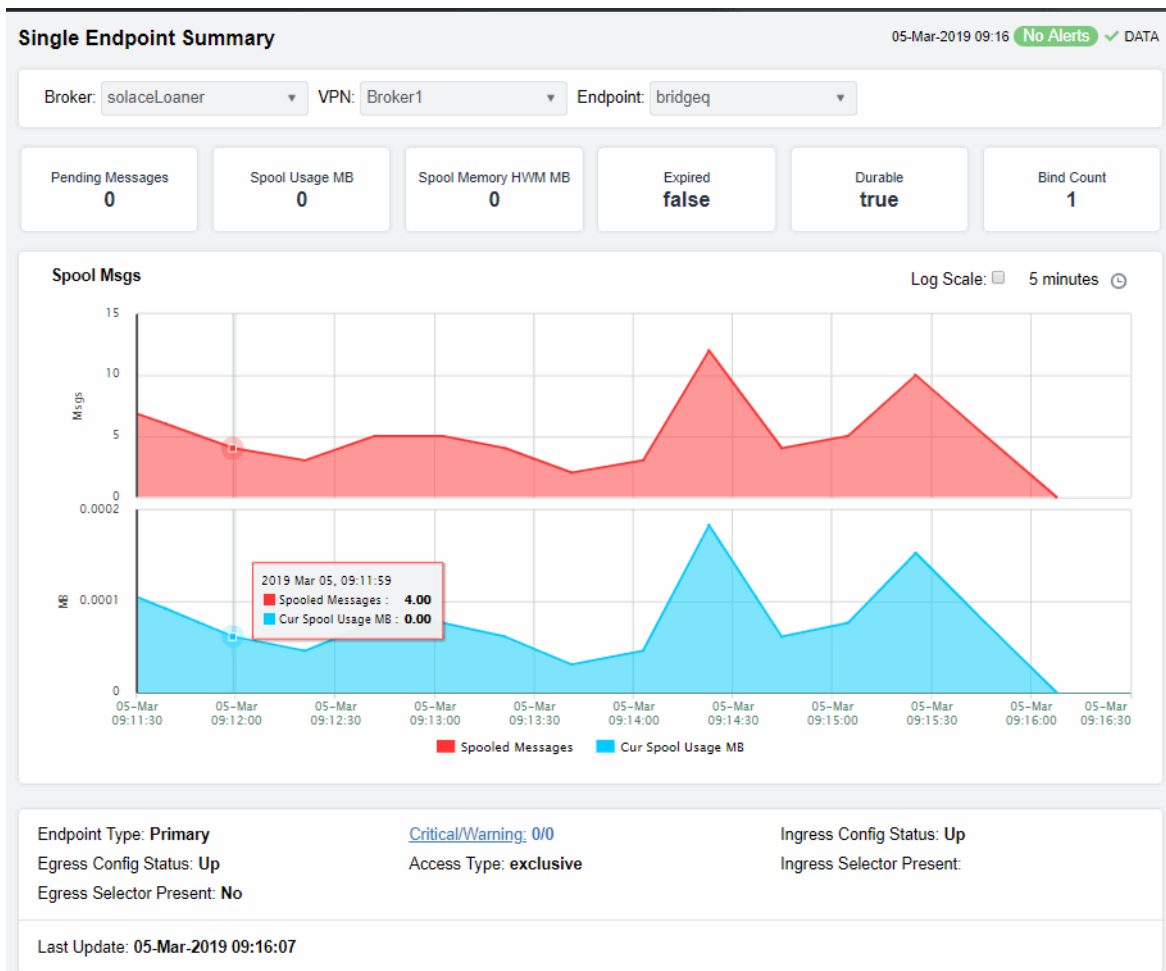
Broker:	Select All brokers or a particular broker.
VPN:	Select All VPNs or a particular VPN.
Show Ingress Config Status Down Only	When checked, only shows endpoints configured on a VPN that currently have this status.
Show:	Select to show All , Expired or Unexpired endpoints configured on the VPN(s).
Type:	Select to show All , Queues Only or Topic Endpoint Only on the VPN(s).
Filter Endpoint Name:	Enter a string to limit list in the table.
Endpoints:	The number of rows currently in the table.
VPN	The name of the VPN.
Endpoint	The name of the endpoint.
Alert Level	<p>The current alert severity in the row.</p> <ul style="list-style-type: none"> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold. Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold. Green indicates that no metrics have exceeded their alert thresholds.
Alert Count	The total number of active alerts for the endpoint.

Bind Count	The total number of binds connected to the endpoint.
Endpoint Type	The type of endpoint (either queue or topic).
Durable	Displays whether or not the endpoint is durable (checked) or non-durable (unchecked). Durable endpoints remain after an broker restart and are automatically restored as part of an broker's backup and restoration process.
In Config Status	Refer to Solace documentation for more information.
Out Config Status	Refer to Solace documentation for more information.
Type	Refer to Solace documentation for more information.
Access Type	Refer to Solace documentation for more information.
Spooled Messages	The total number of spooled messages on the endpoint.
Spool Usage (MB)	The total spool usage consumed on the endpoint (in megabytes).
High Water Mark (MB)	The highest level of spool usage on the endpoint (in megabytes).
In Selector	Refer to Solace documentation for more information.
Out Selector	Refer to Solace documentation for more information.
Expired	When checked, performance data about the endpoint has not been received within the time specified.
Time Stamp	The date and time the row of data was last updated.

Endpoint Summary

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

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



Spooled Messages

The total number of spooled messages on the endpoint.

Spool Usage (MB)

The current spool usage consumed on the endpoint (in megabytes).

Spool Memory HWM MB

Refer to Solace documentation for more information

Expired

When **true**, performance data about the endpoint has not been received within the time specified.

Durable

Displays whether or not the endpoint is durable (checked) or non-durable (unchecked). Durable endpoints remain after an broker restart and are automatically restored as part of an broker's backup and restoration process.

Bind Count

The total number of binds connected to the endpoint.

Trend Graphs

Traces the sum of metrics for the endpoint.

- **Spooled Msgs**: The amount of spooled messages, in megabytes.
- **Cur Spool Usage**: The amount of space used by spooled messages, in megabytes.


Log Scale


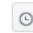
Select to enable a logarithmic scale. Use **Log Scale** to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log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


Base at Zero



Select to use zero (0) as the Y axis minimum for all graph traces.

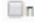
Time Settings

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

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



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

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

Endpoint Type

The type of endpoint.

Egress Config Status

Refer to Solace documentation for more information.

Egress Selector Present

Refer to Solace documentation for more information.

Last Update

The date and time of the last data update.

Critical/Warning

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

Access Type

Refer to Solace documentation for more information.

Ingress Config Status

Refer to Solace documentation for more information.

Ingress Selector Present

Refer to Solace documentation for more information.

Capacity

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

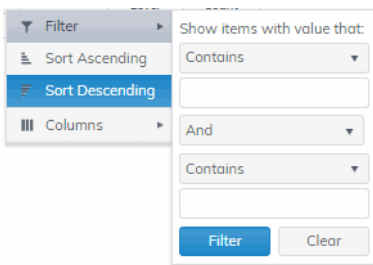
- **"Broker Capacity Table"**: View client, spool usage, incoming messages, outgoing messages, incoming bytes, and outgoing bytes data for all brokers.
- **"Broker Capacity - Summary"**: View client, spool usage, incoming messages, outgoing messages, incoming bytes, and outgoing bytes data for a specific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

Broker Capacity Table

View current and HWM (high water mark for the last 30 days) capacity utilization data for all brokers. By default, a subset of available metrics is shown. Use **More Columns/Less Columns** to toggle to the complete set of metrics available (and back to the subset).




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

Search by clicking the right side of a column heading/**Filter** to open the Search, Sort and Choose Columns dialog:



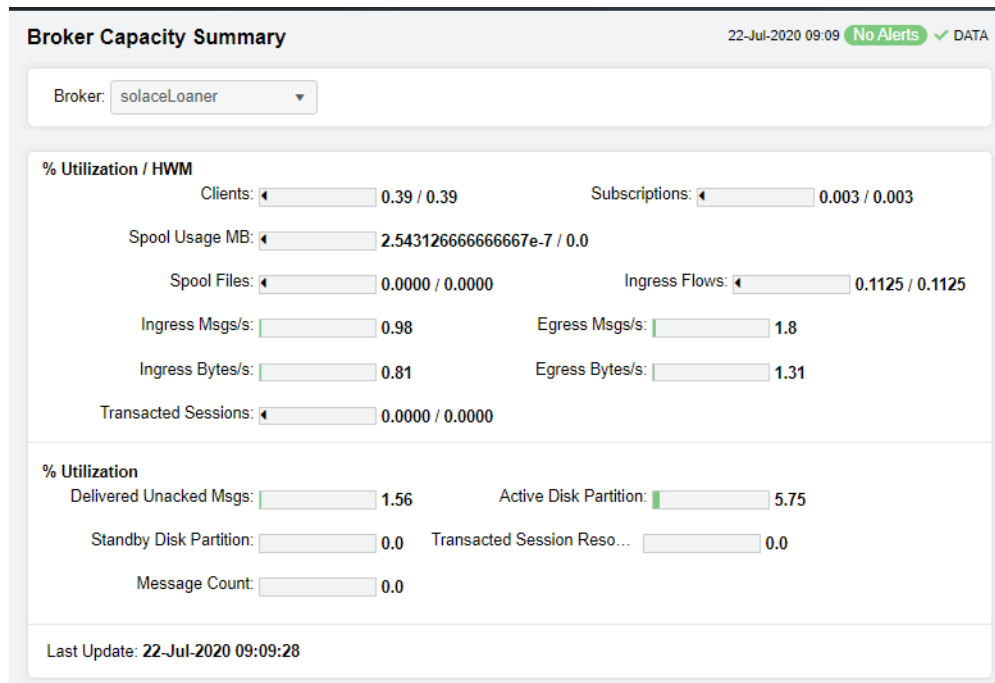
Double-click a row to drill down and investigate in the **"Broker Capacity - Summary"** display.

Broker	Alert Level	Alert Count	Current Client Connections	Connections HWM	Connections Max	Connections Used %	Connect Used HV
VMR_112-SP	✓		2	2	100	2.0	
VMR_114-US-EAST-SP	✓		3	3	100	3.0	
VMR_142-US-EAST-SP	✓		3	3	100	3.0	
VMR_144-SP	✓		4	4	100	4.0	
VMR_209	✓		3	3	100	3.0	

Broker	The name of the broker.
Alert Level	<p>The maximum level of alerts in the row:</p> <p> Red indicates that one or more metrics exceeded their ALARM LEVEL threshold.</p> <p> Yellow indicates that one or more metrics exceeded their WARNING LEVEL threshold.</p> <p> Green indicates that no metrics have exceeded their alert thresholds.</p>
Alert Count	The total number of active alerts.
Current Client Connections	The current number of clients connected.
Connections HWM	The greatest number of connections in the last 30 days.
Connections Max	The greatest number of connections since the broker last started.
Connections Used %	The current amount of connections used, in percent.
Connections Used HWM %	The greatest amount of connections used, in percent, in the last 30 days.
Cur Spool Usage MB	The current amount of used spool disk, in megabytes.
Cur Spool Usage HWM	The greatest amount of spool disk used in the last 30 days.
Spool Size Max MB	The amount of allocated spool disk.
Current Spool Usage %	The current amount of used spool disk, in percent.
Current Spool Usage % HWM	The greatest amount of used spool disk in the last 30 days, in percent.
Delivered Unacked Msgs Util %	Refer to Solace documentation for more information.
Ingress Flow Count	The number of ingress flows.
Ingress Flow HWM	The greatest number of ingress flows in the last 30 days.
Ingress Flows Allowed	The maximum number of ingress flows allowed.
Ingress Flow Count %	The amount of ingress flows in percent.
Ingress Flow Count HWM %	The greatest amount of ingress flows in the last 30 days, in percent.
Ingress Msgs/s	The number of ingress messages per second.
Ingress Msgs/s HWM	The greatest number of ingress messages per second in the last 30 days.
Cur Egress Msgs/s	The number of egress messages per second.
Egress Msgs/s HWM	The greatest number of egress messages per second in the last 30 days.
Egress Bytes/s HWM	The greatest amount of egress, in bytes per second, in the last 30 days, in percent.
Time Stamp	The date and time the row of data was last updated.

Broker Capacity - Summary

This display, a pivoted view of the “[Broker Capacity Table](#)”, allows you to view current and HWM (high water mark for the last 30 days) capacity utilization data for a single broker. Current values are represented by the green bar next to each field. Select a broker from the drop-down menu to view client, spool usage, incoming message, outgoing message, incoming bytes, and outgoing bytes data for the broker. You can mouse over to see rate-per-second values for **Egress** and **Ingress Bytes/s** and **Egress** and **Ingress Msgs/s**.



% Utilization/HWM The following fields show two values: the current % (percent) utilization / peak capacity utilization (HWM) for the last 30 days. HWM values are represented by the black triangular pointer on a bar that reads from 0 to 100%.

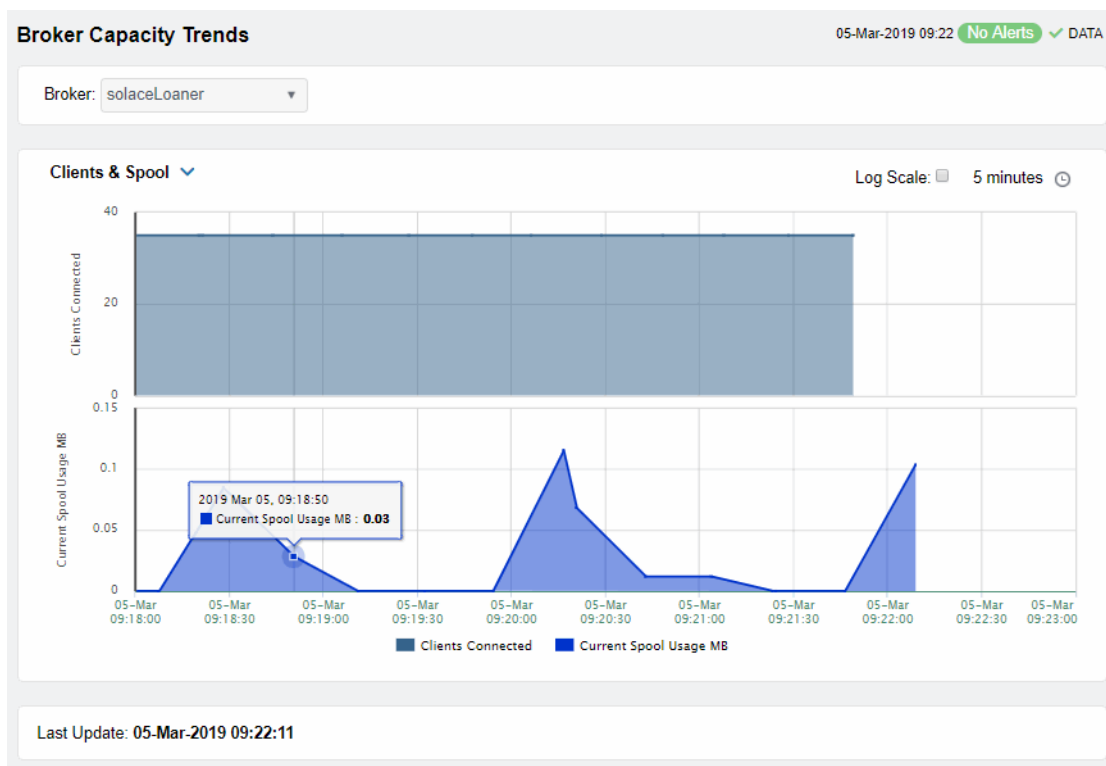
Clients	The clients connected to the broker.
Spool Files	The number of spool files on the broker.
Egress Msgs/s	The outgoing messages per second for the broker.
Transacted Sessions	The number of transacted sessions on the broker.
Subscriptions	The number of subscriptions on the broker.
Ingress Flows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

% Utilization The following fields show current capacity % utilization.

Delivered Unacked Msgs	The current number of delivered messages that were not acknowledged divided by the maximum number of delivered messages that were not acknowledged allowed on the broker.
Transacted Sessions Reso...	The current number of transacted sessions that were resolved on the broker.
Active Disk Partition	The percentage of available active disk partition that is used.
Message Count	The current number of messages on the broker.
Standby Disk Partition	The percentage of available standby disk partition that has been used.
Last Update	The date and time of the last data update.




Broker Capacity Trends


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





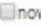
Clients & Spool The trend graph traces the following performance metrics:

- Clients Connected:** The current number of clients connected to the broker.
- Current Spool Usage:** The current spool usage, in megabytes, on the broker.

Message Flow	<p>The trend graph traces the following:</p> <p>Ingress Msgs/sec: The number of incoming messages per second on the broker.</p> <p>Egress Msgs/sec: The number of outgoing messages per second on the broker.</p>
Throughput	<p>The trend graph traces the following:</p> <p>Ingress KB/sec: The amount of incoming per second, in KB, on the broker.</p> <p>Egress KB/sec: The number of outgoing data per second, in KB, on the broker.</p>
Log Scale	<p>Select to enable a logarithmic scale. Use Log Scale to see usage correlations for data with a wide range of values. For example, if a minority of your data is on a scale of tens, and a majority of your data is on a scale of thousands, the minority of your data is typically not visible in non-log scale graphs. Log Scale makes data on both scales visible by applying logarithmic values rather than actual values to the data.</p>
Base at Zero	<p>Select to use zero (0) as the Y axis minimum for all graph traces.</p>
Time Settings	<p>By default, the time range end point is the current time. To change the time range, click the Time Settings  and either:</p> <ul style="list-style-type: none"> choose a Time range from 5 Minutes to 7 Days in the drop-down menu. specify begin/end dates using the calendar . specify begin/end time using the clock .



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

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

Syslog Events

The Solace Syslog Events displays allows you to supervise the last Syslog event messages from the Solace Message Brokers that have been configured for Syslog monitoring. See Solace product documentation for an in depth description of Syslog monitoring in Solace products and how to configure Message Brokers and the Syslog destination.

These displays require the Solace Event Module from the RTView Solace Monitor to be properly configured with a Syslog destination and running. Displays in this View are:

- **"Syslog Events Table"**: A tabular view of all available data for all Syslog events configured on a specific broker.
- **"Syslog Event Summary"**: Expanded/Summary information for a single Syslog event.

Syslog Events Table

This display lists all Syslog events collected from all Solace brokers. Each row in the table is a different message. Use the drop-down menus to filter the list by **Connection**, **Scope** and alert **Severity** level. Filter messages per single broker or all brokers. Click a column header to sort column data in numerical, alphabetical or chronological order. Click a row to investigate the Syslog event in the [“Syslog Event Summary”](#) display.

Search by clicking the right side of a column heading/**Filter** to open the Search, Sort and Choose Columns dialog:

The dialog box is titled 'Filter' and contains the following elements:

- A 'Filter' button with a downward arrow.
- A 'Sort Ascending' button with an upward arrow.
- A 'Sort Descending' button with a downward arrow.
- A 'Columns' button with a rightward arrow.
- A text input field labeled 'Show items with value that:'.
- A dropdown menu with 'Contains' selected.
- A text input field.
- A dropdown menu with 'And' selected.
- A dropdown menu with 'Contains' selected.
- A text input field.
- A 'Filter' button.
- A 'Clear' button.

Solace Syslog Events Table 27-Mar-2019 12:33 ✓ DATA

Connection: - All - More Columns

Scope: - All - Severity: - All - Show: All Events: 3,194

time_stamp	Connection	Solace Scope	Host	Facility	Solace Event...	Syslog Seve...	Type	Message	Additional Fi...	Cl
26-Mar-2019 12:	VMR-47	SYSTEM	ip-172-30-1-144	19	SYSTEM_AUTH	notice	SYSLOG_MSG	Tue Mar 26 19:3	("sessionType":	
26-Mar-2019 12:	VMR-118	SYSTEM	ip-172-30-1-112	19	SYSTEM_AUTH	notice	SYSLOG_MSG	Tue Mar 26 19:3	("sessionType":	
26-Mar-2019 12:	VMR-47	SYSTEM	ip-172-30-1-144	19	SYSTEM_AUTH	notice	SYSLOG_MSG	Tue Mar 26 19:3	("sessionType":	
26-Mar-2019 12:	VMR-118	SYSTEM	ip-172-30-1-112	19	SYSTEM_AUTH	notice	SYSLOG_MSG	Tue Mar 26 19:3	("sessionType":	
26-Mar-2019 12:	VMR-47	SYSTEM	ip-172-30-1-144	19	SYSTEM_AUTH	notice	SYSLOG_MSG	Tue Mar 26 19:3	("sessionType":	
26-Mar-2019 12:	VMR-118	SYSTEM	ip-172-30-1-112	19	SYSTEM_AUTH	notice	SYSLOG_MSG	Tue Mar 26 19:3	("sessionType":	
26-Mar-2019 12:	VMR-47	SYSTEM	ip-172-30-1-144	19	SYSTEM_AUTH	notice	SYSLOG_MSG	Tue Mar 26 19:3	("sessionType":	
26-Mar-2019 12:	VMR-118	SYSTEM	ip-172-30-1-112	19	SYSTEM_AUTH	notice	SYSLOG_MSG	Tue Mar 26 19:3	("sessionType":	
26-Mar-2019 12:	VMR-47	SYSTEM	ip-172-30-1-144	19	SYSTEM_AUTH	notice	SYSLOG_MSG	Tue Mar 26 19:3	("sessionType":	
26-Mar-2019 12:	VMR-118	SYSTEM	ip-172-30-1-112	19	SYSTEM_AUTH	notice	SYSLOG_MSG	Tue Mar 26 19:3	("sessionType":	
26-Mar-2019 12:	VMR-47	SYSTEM	ip-172-30-1-144	19	SYSTEM_AUTH	notice	SYSLOG_MSG	Tue Mar 26 19:3	("sessionType":	
26-Mar-2019 12:	VMR-118	SYSTEM	ip-172-30-1-112	19	SYSTEM_AUTH	notice	SYSLOG_MSG	Tue Mar 26 19:3	("sessionType":	
26-Mar-2019 12:	VMR-47	SYSTEM	ip-172-30-1-144	19	SYSTEM_AUTH	notice	SYSLOG_MSG	Tue Mar 26 19:3	("sessionType":	
26-Mar-2019 12:	VMR-118	SYSTEM	ip-172-30-1-112	19	SYSTEM_AUTH	notice	SYSLOG_MSG	Tue Mar 26 19:3	("sessionType":	

Page 1 of 80 1 - 40 of 3194 items

Connection

Select the connection string assigned when the message brokers connection properties were added with the RTView Configuration Application.

More/Fewer Columns

Switches to another syslog events table display containing the full set of columns coming from Syslog.

Scope:	<p>This drop down selects the type of the event. The SYSTEM events are coming from conditions related to the state of the message broker. VPN events are events with the state of the message brokers VPNs. CLIENT events refer to the state of clients executions in the messaging infrastructure.</p> <p>Available options are:</p> <ul style="list-style-type: none"> • SYSTEM • VPN • CLIENT • ALL shows messages from all sources.
Severity:	<p>Selects the severity level of the events that will be presented in the table. All options go from the less severe to the most important to the health of the systems unless one specifies one single type of severity. For instance, Warning will only show the events that are defined as Warning, filtering out events more damaging, whereas Warning or higher will show all Syslog events that are either Warning, Error, Alert or Emergency. To avoid missing any key event, selection of Warning or higher is recommended.</p> <p>Available options are:</p> <ul style="list-style-type: none"> • INFO • NOTICE • NOTICE or higher • WARN • WARN or higher • ERROR • ERROR or higher • CRITICAL • ALERT • EMERGENCY • ALL shows messages regardless of severity level from all sources.
Show:	<p>Selects the Expiration flag of the event. Due to the large number of events that can exist, it is recommended to select Unexpired Only to see exclusively the events that are active.</p> <p>Available options are:</p> <ul style="list-style-type: none"> • Expired Only • Unexpired Only • ALL shows expired and unexpired messages from all sources.
Events:	The number of events currently shown in the table.
Time Stamp	The date and time the row of data was last updated.

Syslog Event Summary

View status and configuration details for a single Syslog Event, such as the **Syslog Severity Code**, the **Intended Support Group**, and whether the event is **Clearable** and **Expired**.

Select a **Connection** and an **ID** from the drop-down menus.

Refer to Solace product documentation for additional information about values provided in this display.

Solace Syslog Event Summary
16-Jun-2020 14:48
No Alerts
DATA

Connection: centos7-1

ID: vm-centos7-1_SYSTEM_AD_MSG_CACHE_USAGE_HIGH_CLEAR_1592336901637

Connection: centos7-1	Solace Scope: SYSTEM	Host: vm-centos7-1
Router: vm-centos7-1	Facility: 19	

Solace Event ID: **SYSTEM_AD_MSG_CACHE_USAGE_HIGH_CLEAR**

ID: **vm-centos7-1_SYSTEM_AD_MSG_CACHE_USAGE_HIGH_CLEAR_1592336901637**

Solace Event Source: event.log	Syslog Severity: Informational	Clearable: false
Is Clear Event: true	Syslog Severity Code: 6	Type: SYSLOG_MSG

Message:
Tue Jun 16 10:54:16 PDT 2020 vm-centos7-1 event: SYSTEM: SYSTEM_AD_MSG_CACHE_USAGE_HIGH_CLEAR: -- Guaranteed Message Cache Usage Threshold 6.0 (60%) cleared: 5.2

Solace Description: **Guaranteed Message Cache Usage Threshold 6.0 (60%) cleared: 5.2**

Additional Fields: {"actual": 5,"threshold": 6,"thresholdPct": 60,"actualTenths": 2,"thresholdTenths": 0}

Clearable: false	Is Clear Event: true	Intended Support Group: APPLICATION
Solace Vpn: -	Solace Client: -	Solace Company: -
Solace Application: -	Solace Environment: -	Expired: false

Last Update: **16-Jun-2020 14:48:21**

Connection:	Select a connection.
ID:	Select an event ID.
Connection	The name of the selected connection.
Router	The name of the router for the selected event.
Solace Scope	The type of Syslog event: <ul style="list-style-type: none"> • SYSTEM events are coming from conditions related to the state of the message broker. • VPN events are events with the state of the message brokers VPNs. • CLIENT events refer to the state of clients executions in the messaging infrastructure.
Facility	The generating facility for the event. Event broker Syslog messages belong to a facility, which is a group of messages that are either generated by the same software process, or concern a similar event broker subsystem condition or activity (such as debugging attempts).
Host	The name of the host associated with the event.
Solace Event ID	The type of Solace event.
ID	The unique identifier for the event.

Solace Event Source	The name of the log file in which the event is recorded.
Is Clear Event	true/false
Syslog Severity:	The Syslog severity level for the event: 0 - Emergency: system is unusable. 1 - Alert: action must be taken immediately. 2 - Critical: critical conditions. 3 - Error: error conditions. 4 - Warning: warning conditions. 5 - Notice: normal but significant condition. 6 - Informational: informational messages. 7 - Debug: debug-level messages.
Syslog Severity Code:	0 - Emergency: system is unusable. 1 - Alert: action must be taken immediately. 2 - Critical: critical conditions. 3 - Error: error conditions. 4 - Warning: warning conditions. 5 - Notice: normal but significant condition. 6 - Informational: informational messages. 7 - Debug: debug-level messages.
Clearable:	true/false
Type:	The type of the selected event.
Message:	The Syslog message content.
Solace Description	The Message Text for this event.
Additional Fields	Additional information associated with the event (formatted as Name:Value pairs in a JSON object).
Clearable	true/false
Solace Vpn	The Solace VPN identifier associated with this event.
Solace Application	The Solace Application identifier associated with this event.
Is Clear Event	true/false
Solace Client	The Solace Client identifier associated with this event.
Solace Environment	The Solace environment identifier associated with this event .
Intended Support Group	The intended support group for this event.
Solace Company	The Solace Company identifier associated with this event.
Expired	true/false
Last Update	The date and time of the last data update.

Drill Down Displays

The displays described in this section are only accessible from other displays. These displays are used for managing alerts at the component level.

This View includes the following displays:

- [“Alerts History Table - HTML”](#): Track historical alerts that have occurred in your monitoring system.
- [“Alerts Table by Component - HTML”](#): Track alerts associated with CIs shown in a display.
- [“Alert Detail for Component - HTML”](#): Investigate an alert instance and its history.
- [“Alert Configuration for Component - HTML”](#): Refine alert threshold settings.

Alerts History Table - HTML

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

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

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

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

You can search, filter, sort and choose columns to include by clicking a column header icon (to the right of each column label) and selecting **Filter**, **Sort Ascending**, **Sort Descending** or **Columns**. Right-click on a table cell to **Export to Excel**.

Alerts History Table 07-Oct-2019 12:12 DATA

Data Server: RTV-DATA-SOLMON Expand Alert Index: ☐ History Alerts: ☒

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

Alert Level	Ack	Cleared	Alert Name	Alert Index	Alert Text	Owner	Id	Source	Row Update Time
			SolMsgRouterPendingMsgsHig	VMR-209	High Alert Limit exceeded		120917		2019-Oct-07 11:39:...
			SolMsgRouterPendingMsgsHig	VMR-114-US-EAST	High Alert Limit exceeded		120918		2019-Oct-07 11:39:...
		✓	SolMsgRouterPendingMsgsHig	solDemo	High Warning Limit exce		120916		2019-Oct-07 11:39:...
			SolMsgRouterPendingMsgsHig	VMR-144	High Alert Limit exceeded		120919		2019-Oct-07 11:39:...
			SolMsgRouterPendingMsgsHig	VMR-112	High Alert Limit exceeded		120920		2019-Oct-07 11:39:...
		✓	SolMsgRouterPendingMsgsHig	VMR-209	High Alert Limit exceeded		120917		2019-Oct-07 11:39:...
		✓	SolMsgRouterPendingMsgsHig	VMR-114-US-EAST	High Alert Limit exceeded		120918		2019-Oct-07 11:39:...
		✓	SolMsgRouterPendingMsgsHig	VMR-144	High Alert Limit exceeded		120919		2019-Oct-07 11:40:...
		✓	SolMsgRouterPendingMsgsHig	VMR-112	High Alert Limit exceeded		120920		2019-Oct-07 11:40:...
			SolMsgRouterPendingMsgsHig	solDemo	High Alert Limit exceeded		120921		2019-Oct-07 11:40:...
		✓	SolMsgRouterPendingMsgsHig	solDemo	High Alert Limit exceeded		120921		2019-Oct-07 11:49:...
			SolMsgRouterPendingMsgsHig	VMR-209	High Alert Limit exceeded		120922		2019-Oct-07 11:50:...
			SolMsgRouterPendingMsgsHig	VMR-114-US-EAST	High Alert Limit exceeded		120923		2019-Oct-07 11:50:...
			SolMsgRouterPendingMsgsHig	VMR-144	High Alert Limit exceeded		120924		2019-Oct-07 11:50:...
			SolMsgRouterPendingMsgsHig	VMR-112	High Alert Limit exceeded		120925		2019-Oct-07 11:50:...
		✓	SolMsgRouterPendingMsgsHig	VMR-209	High Alert Limit exceeded		120922		2019-Oct-07 11:50:...
		✓	SolMsgRouterPendingMsgsHig	VMR-114-US-EAST	High Alert Limit exceeded		120923		2019-Oct-07 11:50:...
		✓	SolMsgRouterPendingMsgsHig	VMR-144	High Alert Limit exceeded		120924		2019-Oct-07 11:51:...
		✓	SolMsgRouterPendingMsgsHig	VMR-112	High Alert Limit exceeded		120925		2019-Oct-07 11:51:...
			SolMsgRouterPendingMsgsHig	solDemo	High Alert Limit exceeded		120926		2019-Oct-07 11:51:...

Alerts Table by Component - HTML

As an alternative to the **Alerts Table**, use the **Alerts Table by Component** to track and manage all alerts that are specifically associated with the CIs shown in a display.

You access the **Alerts Table by Component** by clicking (the alert status icon) in the title bar of other displays. The display in which you click is the source display.

Package provides the technology label associated with the alerts shown. For example, **Jvm**, **Tomcat** and **Host** are the technology labels for Java Virtual Machines, Tomcat applications and servers (respectively). These labels are also correlated with the RTView solution package names (for example, the Solution Package for Host Agent). **Category** lists all alert categories related to the source display.

Use the **ACK** and **Cleared** drop-downs to filter the table by **All**, **True** or **False**.

See the **Alert Level** column icon, where:



The alert reached its ALARM LEVEL threshold in the table row.



The alert reached its WARNING LEVEL threshold in the table row.

To investigate, click:

to open the **Alert Detail for Component** where you can see the current and historical conditions that precipitated the alert being executed.

to open the summary display for the CI associated with the alert where you can investigate utilization metrics for the CI leading up to the alert being executed.

You can search, filter, sort and choose columns to include by clicking a column header icon (to the right of each column label) and selecting **Filter**, **Sort Ascending**, **Sort Descending** or **Columns**. Right-click on a table cell to **Export to Excel**. Use **Ctrl** + click or **Shift** + click to select multiple alerts.

With one or more alerts selected, click **Own** to set the alert(s) owner field, **Acknowledge** to acknowledge the alert(s), **Unacknowledge** to clear the acknowledgement on previously acknowledged alert(s), **Add Comment** to add a comment to the alert(s).

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

Alerts Table by Component 02-May-2019 11:05:09 DATA OK

Package: Host Category: CPU:Network:Storage Cleared: False ACK: False

Alert Count: 16

Row	Update Time	Acknowledge	Cleared	Alert Level	Alert Name	Alert Index Values	
2018-Nov-09 23:54:0				🔔	HostCpuPercentHigh	SL-DEMO:SLHOST16(sl_qa)	High V
2018-Oct-01 06:20:10				⚠️	HostCpuPercentHigh	SL-DEMO:SLHOST17(sl_amx)	High A
2019-May-02 03:28:5				🔔	HostMemoryUsedHigh	SL-DEMO-LX;192.168.200.92	High V
2018-Oct-01 06:19:36				⚠️	HostVirtualMemoryUsedH	SL-DEMO:SLHOST17(sl_amx)	High A
2018-Oct-01 06:18:36				🔔	HostMemoryUsedHigh	SL-DEMO:SLHOST17(sl_amx)	High V
2018-Jan-12 11:38:56				⚠️	HostCpuPercentHigh	SL-DEMO-LX;192.168.200.205	High A
2019-May-02 10:40:3				⚠️	HostVirtualMemoryUsedH	SL-DEMO-LX;192.168.200.42	High A
2019-Apr-25 10:19:43				🔔	HostMemoryUsedHigh	SL-DEMO:SLHOST8	High V
2018-Jun-19 09:22:23				⚠️	HostCpuPercentHigh	SL-DEMO-LX;192.168.200.202	High A
2018-Nov-09 10:33:54				⚠️	HostVirtualMemoryUsedH	SL-DEMO:SLHOST16(sl_qa)	High A
2018-May-01 22:45:4				⚠️	HostCpuPercentHigh	SL-DEMO-LX;192.168.200.202	High A

Alert Detail Go to CI Own Acknowledge Unacknowledge

Add Comment Clear All Comments

Alert Detail for Component - HTML

Use the **Alert Detail for Component** display to investigate current and historical activity of a specific alert instance as it applies to the associated CI, and also compare against **Metric History** trends of the associated CI. A trend graph for the CI associated with the alert instance. You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.

Access the **Alert Detail for Component** display by clicking **Details** in the **Alerts Table** or **Alert Detail** in the **Alerts Table by Component** display.

The **Alert History** table at the bottom of the display contains a row of data for each time the alert instance was updated. See the alert **ID**, **Row Update Time**, **Cleared** status and **Reason**, **Owner** and the **Alert Level** column icon, where:



The alert reached its ALARM LEVEL threshold in the table row.




The alert reached its WARNING LEVEL threshold in the table row.

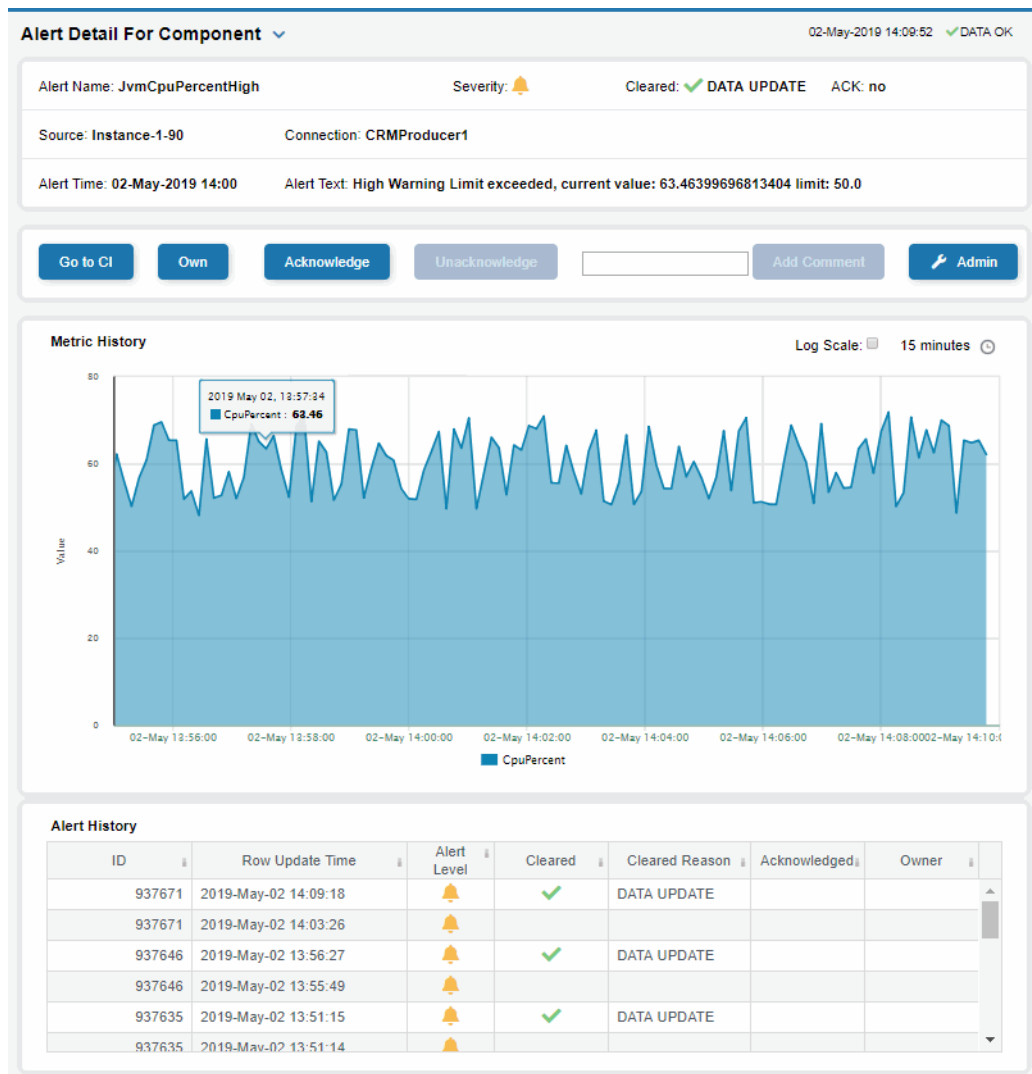
You can search, filter, sort and choose columns to include by clicking a column header icon (to the right of each column label) and selecting **Filter**, **Sort Ascending**, **Sort Descending** or **Columns**. Right-click on a table cell to **Export to Excel**. Use **Ctrl** + click or **Shift** + click to select multiple alerts.

To investigate, click:

[Go to CI](#)

to see utilization conditions for the CI associated with the alert in a summary display.

 **Admin** to open the **Alert Configuration for Component** display where you can see, modify and refine alert threshold settings for that particular alert. A trend graph traces the relevant alert metric for the CI so you can adjust thresholds in real-time.



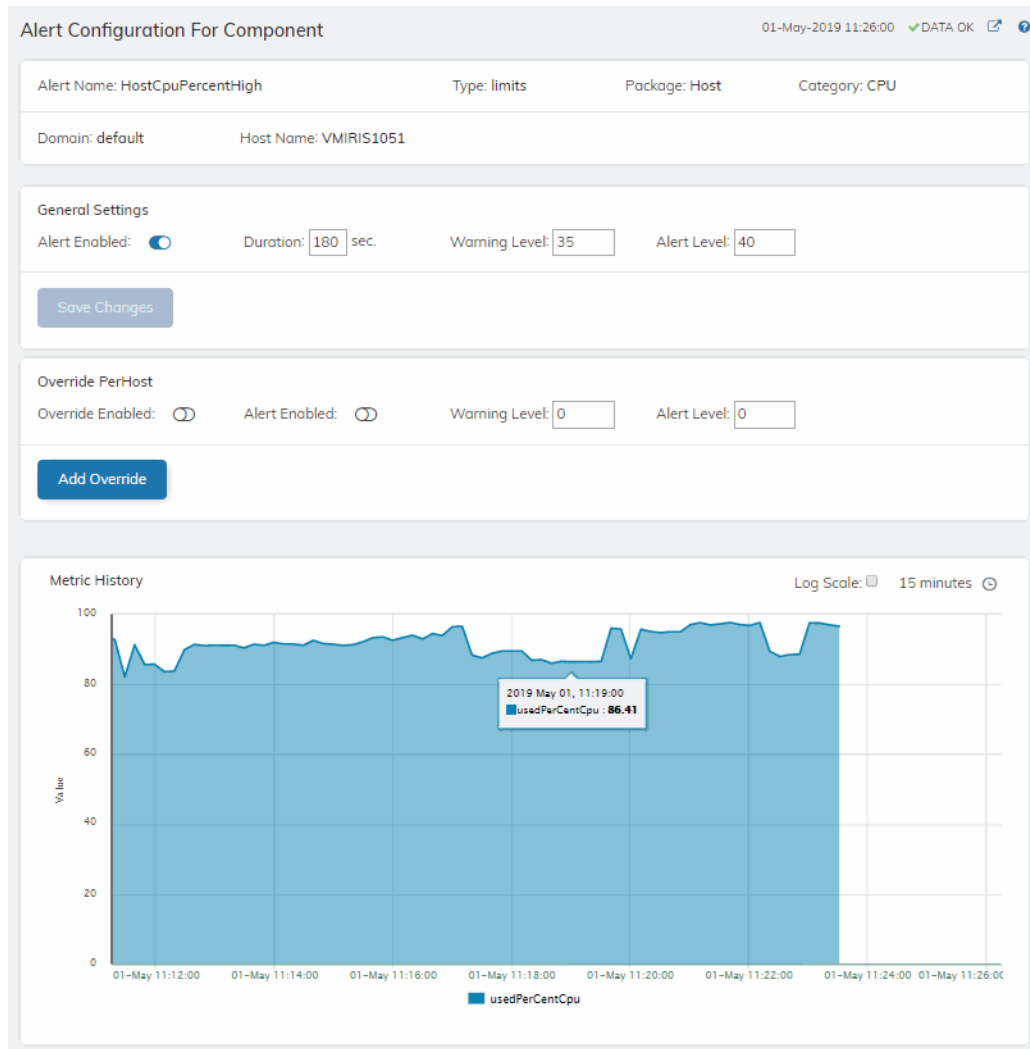
Alert Configuration for Component - HTML

Use the **Alert Configuration for Component** display to see, modify and refine alert threshold settings for a particular alert. A trend graph traces the history of the relevant metric for this alert so you can adjust thresholds in real-time. You can also modify alert thresholds, add an override alert and toggle ON or OFF  both global and override alerts.

Access the **Alert Configuration for Component** display by clicking  **Admin** in the **Alert Detail for Component** display.

The bottom half of the display provides a **Metric History** trend graph which traces the performance metric pertaining to the alert. You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.

You must be logged in as rtvalertmgr or rtvadmin to modify alerts.



Alerts

This section describes the “[Alerts Table](#)” and how to use it to manage open alerts.

Alerts Table

Use this display to track and manage all alerts that have occurred in the system, where:



One or more alerts exceeded their ALARM LEVEL threshold in the table row



One or more alerts exceeded their WARNING LEVEL threshold in the table row

You can search, filter, sort and choose columns to include by clicking a column header icon (located to the right of each column label) and selecting **Filter**, **Sort Ascending**, **Sort Descending** or **Columns**. Use the **Ack'd** and **Cleared** drop-downs to filter the table by those columns. Right-click on a table cell to **Export to Excel** or **Copy Cell Value**. Use **Ctrl** + click or **Shift** + arrow to select multiple alerts. To investigate, select one alert and click:

Details

to open the **Component Alert Detail** display to get details about that particular alert instance as it specifically applies to the associated CI.

CI

to see utilization conditions for the CI associated with the alert during the seconds (minutes, hours or days) leading up to the alert being executed in a summary display.

With one or more alerts selected, you can click **Own** to set the alert(s) owner field, **Ack** to acknowledge the alert(s), **Unack** to clear the acknowledgement on previously acknowledged alert(s), **Clear** to set the **Cleared** flag on the selected alert(s), **Comment** to add a comment to the alert(s) and **CI** to get details about the CI associated with the alert (these buttons are enabled when you click one or more alerts).

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

Alerts Table											
30-Apr-2019 13:47:46 ✓ DATA											
<div> <div>Own</div> <div>Ack</div> <div>Unack</div> <div>Clear</div> <div>Comment</div> <div>Details</div> <div>CI</div> </div> <div> <div>Ack'd: all</div> <div>Cleared: false</div> <div>Cmdb Filter: *****</div> <div>Alert Count: 92</div> </div>											
Time	Ack	Clr	Sev	Alert Name	Alert Text	Owner	ID	Source	Comments	CI	
2019-Apr-30 00:04:07			⚠	JvmNotConnected	Server disconnected		1043	RTV-DATA-TIB		win4	
2019-Apr-30 01:34:49			⚠	JvmNotConnected	Server disconnected		1009	Z-SIMDATA-1		local	
2019-Apr-30 01:34:49			⚠	JvmNotConnected	Server disconnected		1008	Z-SIMDATA-1		local	
2019-Apr-30 01:34:49			⚠	JvmNotConnected	Server disconnected		1007	Z-SIMDATA-1		local	
2019-Apr-30 01:34:49			⚠	JvmNotConnected	Server disconnected		1006	Z-SIMDATA-1		local	
2019-Apr-30 01:34:49			⚠	JvmNotConnected	Server disconnected		1005	Z-SIMDATA-1		local	
2019-Apr-30 01:34:49			⚠	JvmNotConnected	Server disconnected		1004	Z-SIMDATA-1		local	
2019-Apr-30 01:34:49			⚠	JvmNotConnected	Server disconnected		1003	Z-SIMDATA-1		local	
2019-Apr-30 01:34:49			⚠	JvmNotConnected	Server disconnected		1002	Z-SIMDATA-1		local	
2019-Apr-30 01:34:49			⚠	JvmNotConnected	Server disconnected		1001	Z-SIMDATA-1		local	
2019-Apr-30 01:34:49			⚠	JvmNotConnected	Server disconnected		1000	Z-SIMDATA-1		local	
2019-Apr-30 12:01:02			⚠	JvmCpuPercentHigh	High Alert Limit exceed		1064	Z-SIMDATA-1		local	
2019-Apr-30 13:44:01			🔔	JvmCpuPercentHigh	High Warning Limit exc		928739	RTV-DATA-KAF		Insta	
2019-Apr-30 13:47:04			🔔	JvmCpuPercentHigh	High Warning Limit exc		928747	RTV-DATA-KAF		Insta	
2019-Apr-30 01:36:49			🔔	HostCpuPercentHigh	High Warning Limit exc		1010	Z-SIMDATA-1		defa	
2019-Apr-30 01:36:49			🔔	HostCpuPercentHigh	High Warning Limit exc		1010	Z-SIMDATA-1		defa	
2019-Apr-30 02:05:10			⚠	HostCpuPercentHigh	High Alert Limit exceed		1011	Z-SIMDATA-1		defa	

Admin

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

- **"Alert Administration"**: Displays active alerts and provides interface to modify, enable and manage alerts.
- **"Alert Overrides Admin"**: Set and modify alert overrides. Access this display from the **Alert Administration** display.
- **"Cache Table"**: View cached data that RTView is capturing and maintaining, and use this data use this for debugging with SL Technical Support.

Alert Administration

The **Alert Administration** display allows administrators to enable/disable alerts and manage alert thresholds. The table describes the global settings for all alerts on the system.

You can set the **Delay** time (the number of seconds that must pass before an alert is triggered, where **0** sets it to immediately execute).

You can set the **Warning Level** which executes a single warning alert when the number of seconds specified here is exceeded. To set the warning to occur sooner, reduce the **Warning Level** value. To set the warning to occur later, increase the **Warning Level** value.

You can set the **Alarm Level** which executes a single alarm alert when the number of seconds specified here is exceeded. To set the alarm to occur sooner, reduce the **Alarm Level** value. To set the alarm to occur later, increase the **Alarm Level** value.

Note: For low value-based alerts (an alert that executes based on a value going below a certain threshold), to set the alarm to occur sooner you increase the **Alarm Level** value. To set the alarm to occur later, reduce the **Alarm Level** value.

You can apply alert thresholds globally or as an *override*. Setting override alerts allows you to set thresholds for a subset of your resources, or for a single resource (for example, a single server). Override alerts are useful if the majority of your resources require the same threshold setting, but there are a few resources that require a different threshold setting. For example, you might not usually be concerned with execution time at a process level, but perhaps certain processes are critical. In this case, you can apply alert thresholds to each process individually. See below for instructions.

You can filter, sort and choose columns to include by clicking a column header icon (located to the right of each column label) and selecting **Filter**, **Sort Ascending**, **Sort Descending** or **Columns**. Use the **Ack'd** and **Cleared** drop-downs to filter the table by those columns. Right-click on a table cell to **Export to Excel**.

To set thresholds and enable a global alert:

Select an alert and, under **Settings for alert** (in the lower portion of the screen), modify settings for the alert **Delay**, **Warning Level** and/or **Alarm Level** and **Save Settings**. With that alert selected, check the **Alert Enabled** box under **Settings for alert** (in the lower portion of the screen) and **Save Settings**. The **Alert Enabled** box (next to the selected alert) is now checked.

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

To set thresholds and enable an override alert:

To set an override alert, select an alert and click **Override Settings** to open the **Alert Overrides Admin** display.

Alerts Administration 30-Apr-2019 10:34:01 ✓ DATA OK

Package: All http://rtvdemos.sl.com/emdemo_central_rtvquery

Alert Name	Alert Enabled	Alert Delay	Warning Level	Alert Level	Override Count
HostNetworkTxRateHigh	<input type="checkbox"/>	30	50	75	0
HostProcessCountLow	<input type="checkbox"/>	30	15	5	0
HostStateData	<input type="checkbox"/>	30			0
HostStorageUsedHigh	<input type="checkbox"/>	30	80	90	0
HostSwapUsedHigh	<input type="checkbox"/>	30	75	90	0
HostVirtualMemoryUsedHigh	<input type="checkbox"/>	30	75	90	0
JvmCpuPercentHigh	<input checked="" type="checkbox"/>	60	50	70	0
JvmGcDutyCycleHigh	<input type="checkbox"/>	30	50	75	0
JvmMemoryUsedAfterGCHigh	<input type="checkbox"/>	0	1	80	0
JvmMemoryUsedHigh	<input checked="" type="checkbox"/>	60	75	86	0
JvmNotConnected	<input checked="" type="checkbox"/>	60			0
JvmStateData	<input type="checkbox"/>	30			0
JvmThreadCountHigh	<input checked="" type="checkbox"/>	60	8000	12000	0

Page 2 of 5 101 - 200 of 432 items

Settings for alert

Alert Enabled: ☐ Delay: Warning Level: Alert Level:

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

For additional details, see ["Alert Overrides Admin"](#).

Alert Name	The name of the alert.
Alert Enabled	When checked, the alert is enabled globally.
Alert Delay	The amount of time (in seconds) that the value must be above the specified Warning Level or Alarm Level threshold before an alert is executed. 0 is for immediate execution.

Warning Level	The global warning threshold for the selected alert. When the specified value is exceeded a warning is executed.
Alert Level	The global alarm threshold for the selected alert. When the specified value is exceeded an alarm is executed.
Override Count	The number of times thresholds for this alert have been defined individually in the Tabular Alert Administration display. A value of: -0 indicates that no overrides are applied to the alert. -1 indicates that the alert does not support overrides.

Settings for alert

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

Alert Enabled	Check / uncheck this box to enable or disable the selected alert globally.
Delay	Enter the amount of time (in seconds) that the value must be above the specified Warning Level or Alarm Level threshold before the selected alert is executed. 0 is for immediate execution.
Warning Level	Enter the global warning threshold for the selected alert. When the specified value is exceeded a warning is executed. To set the warning to occur sooner, reduce the Warning Level value. To set the warning to occur later, increase the Warning Level value.
Alert Level	Enter the global alarm threshold for the selected alert. When the specified value is exceeded an alarm is executed. To set the alarm to occur sooner, reduce the Alarm Level value. To set the warning to occur later, increase the Alarm Level value. NOTE: For low value-based alerts (such as EmsQueuesConsumerCountLow), to set the alarm to occur sooner, increase the Alarm Level value. To set the alarm to occur later, reduce the Alarm Level value.
Save Settings	Click to apply alert settings for the selected alert.
Original Defaults	Click to revert to original alert settings for the selected alert.
Override Settings	Click to set an alert override in the Alert Overrides Admin display on the selected alert.

Alert Overrides Admin

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

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

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

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

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

You can filter, sort and choose columns to include by clicking a column header icon (located to the right of each column label) and selecting **Filter**, **Sort Ascending**, **Sort Descending** or **Columns**. Use the **Display** drop-down to filter the table to show **All** resources, only resources with the **Overridden** alert applied or **Free** resources (to show only resources without the alert override applied). Right-click on a table cell to **Export to Excel** or **Copy Cell Value**.

To set overrides:

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

To alter overrides:

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

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

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

Alert Overrides Administration 10-Jun-2020 16:25:07 ✓ DATA OK

Alert: SolBrokerEgressByteRateHigh Override Type: PerMsgRouter Display: All

Search: * RegEx: ☐

Connection	Override Enabled	Alert Enabled	Alert Delay	Warning Level	Alert Level
solaceLoaner					
VMR_144-SP					
VMR_209					
VMR_142-US-EAST					
VMR_114-US-EAST					
VMR_112-SP					

Settings for selected index

Override Enabled: ☒ Alert Enabled: ☒ Alert Delay: 30 Warning Level: 8000000

Alert Level: 10000000

Add Override Save Settings Remove Override

Cache Table

View the raw data that RTView is capturing and maintaining to investigate utilization and capacity metrics, as well as connection details, for caches on a data server.

Select a **Data Server** from the drop-down menu. The upper table contains a row of data for each cache on the selected data server. You can see the current number of **Rows** and **Columns** in each table and the amount of **Memory** used. You can also find out the cache **Table** type of which there are five:

- **current** tables show the most recently received values for each index.
- **current_condensed** tables are current tables with primary compaction configured.
- **history** tables show the historical values for each index.
- **history_condensed** tables are history tables with primary compaction configured.
- **history_combo** tables are history tables with primary compaction configured, and which is also configured to store rows of recent raw data followed by rows of older condensed data.

Select a cache to see connection utilization details for that cache in the lower table. The lower table shows the contents of the selected cache table. Available columns vary by cache. For example, a JVM cache table might provide **BootClassPath** and **InputArgument** columns, and a Tomcat cache might provide **RateAccess** and **cacheMaxSize** columns.

You can search, filter, sort and choose columns to include by clicking a column header icon (to the right of each column label) and selecting **Filter**, **Sort Ascending**, **Sort Descending** or **Columns**. Or just click a column header to sort.

Right-click on a table cell to **Export to Excel** or **Copy Cell Value**. Use **Ctrl + click** or **Shift + click** to select multiple alerts. Use **History Tables** to include / exclude history tables in the table. Right-click on a table cell to **Export to Excel** or **Copy Cell Value**.

This low-level option can be useful to identify the source of the problem when the displays are not showing the expected data. Use this data for debugging and troubleshooting with Technical Support.

Cache Table 07-May-2019 14:11 ✓ DATA

Data Server: central-alert History Tables: ☐

Data Server URL: https://rtvdemos.sl.com/emdemo_central_rtvquery

Cache	Table	Rows	Columns	Memory
JmxStatsTotals	current	1	4	441
RtvAlertGroupMap	current	493	3	67424
RtvAlertMapByCI	current	62	5	13614
RtvAlertSourceStats	current	8	2	940
RtvAlertStatsByArea	current	8	9	2930
RtvAlertStatsByAreaAndAlertGroup	current	8	10	3454
RtvAlertStatsByCI	current	59	5	9228
RtvAlertStatsByCIAndAlertGroup	current	59	8	12506

Cache: RtvAlertStatsByCIAndAlertGroup Table: current

time_stamp	CITYPE	CINAME	ALERTGROUP	MaxSeverity	AlertCount
2019-May-07 14:11:33	JVM	localhost:SOLMON_TOM	None	2	1
2019-May-07 14:11:33	JVM	localhost:EMSMON_TOM	None	2	1
2019-May-07 14:11:33	JVM	localhost:EMSMON_DAT	None	2	1
2019-May-07 14:11:33	JVM	localhost:SOLMON_DISF	None	2	1
2019-May-07 14:11:33	JVM	localhost:SOLMON_DAT	None	2	1
2019-May-07 14:11:33	JVM	localhost:EMSMON_DISI	None	2	1
2019-May-07 14:11:33	JVM	localhost:SOLMON_TOM	None	2	1
2019-May-07 14:11:33	JVM	localhost:EMSMON_DAT	None	2	1
2019-May-07 14:11:33	JVM	Instance-1-90;CRMBroke	None	1	1
2019-May-07 14:11:33	JVM	Instance-1-90;CRMZooki	None	1	1
2019-May-07 14:11:33	JVM	Instance-1-171;CRMCon	None	1	1
2019-May-07 14:11:33	JVM	Instance-1-171;CRMCon	None	1	1
2019-May-07 14:11:33	JVM	Instance-1-171;CRMBrok	None	1	1
2019-May-07 14:11:33	JVM	localhost:TMolecule5_2	None	1	1
2019-May-07 14:11:33	JVM	localhost:PMolecule12_1	None	1	1

Page 1 of 2 1 - 40 of 59 items

CHAPTER 7 RTView Manager

Use the RTView Manager application to track the health of your Solace PubSub+ Monitor system: RTView and Tomcat processes, the historian and the Solace data server. RTView Manager runs as a process, separately from your Solace PubSub+ Monitor system, has its own URL / console, as well as its own data server, database, alert notification system and historian.

This chapter contains:

- [“Displays”](#): Describes [“Tomcat Displays”](#), [“JVM Processes Displays”](#), [“RTView Servers Displays”](#) and [“Drilldowns’ Displays”](#) which are accessed from the navigation tree (left panel).
- [“Alerts Displays”](#): Describes the display that is used to manage alerts.
- [“Admin Displays”](#): Describes displays that administrators use to manage alert thresholds.
- [“Modify RTView Manager Settings”](#): Describes how to change default settings for RTView Manager.
- [“Troubleshoot”](#): Tips for resolving technical issues.
- [“Configure Alert Notification”](#): Describes how to setup alert notification.
- [“Configure High Availability”](#): Describes how to configure HA.

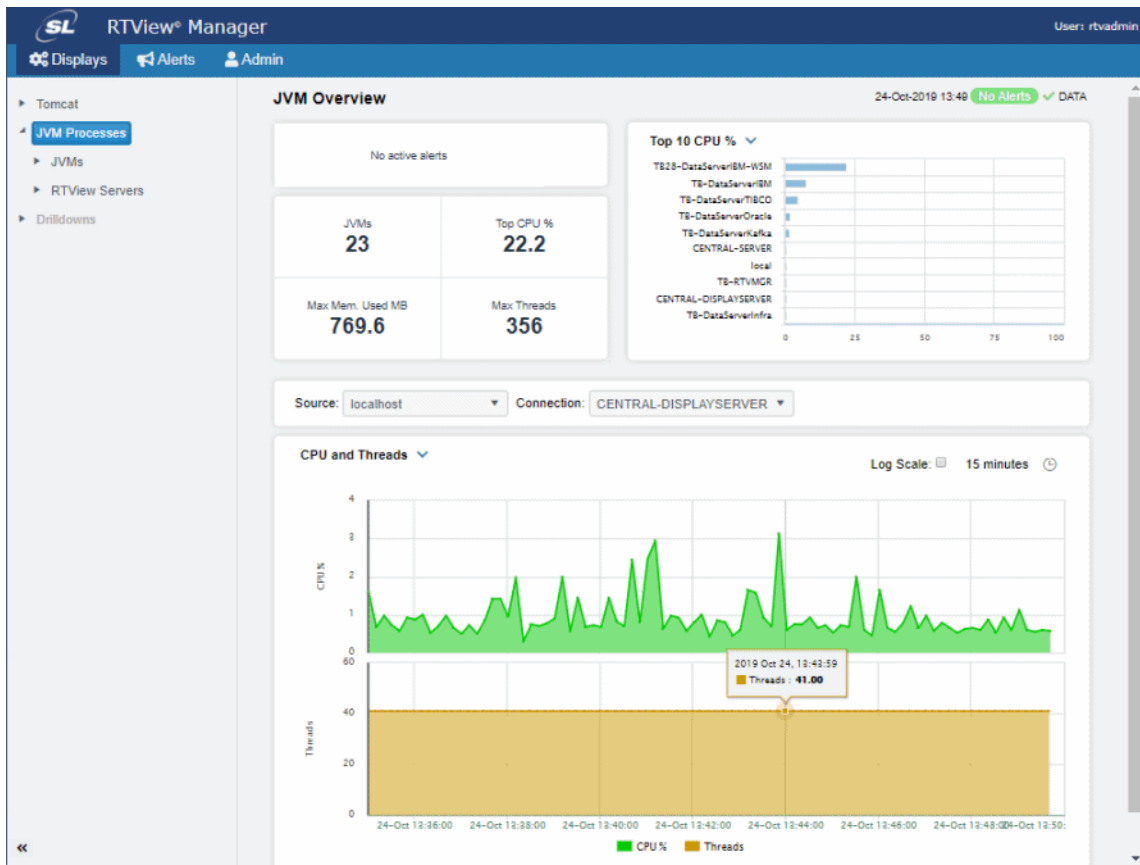
Login to RTView Manager

To access RTView Manager, start the Solace PubSub+ Monitor (if not currently running), then browse to one of the following URLs and login (username/password are rtvadmin/rtvadmin):

http://<ip_address>:3070/rtview-manager if you are running on Jetty.

http://localhost:8068/rtview-manager if you are using Tomcat.

The RTView Manager main console opens.



Displays

The displays that come with RTView Manager are organized by the following Views in the navigation tree:

- [“Tomcat Displays”](#): For monitoring the health of Tomcat servers, applications and all installed web modules. Performance data provided includes current and historic metrics, number of sessions, request rates, cache hit rates and data transmission metrics.
- [“JVM Processes Displays”](#): For monitoring the health of Java Virtual Machine (JVM) processes. JVM metrics track garbage collection information and trends, including memory usage before and after garbage collection, duration and duty cycles. This, combined with tracking of JVM memory pool trends, enables you to be notified of memory leaks, unusual garbage collection activities and CPU and memory resource issues automatically with minimal user analysis, speeding the discovery of the root cause of any issue. It also monitors a Java Virtual Machine’s memory heap, non-heap memory, monitor threads and other key metrics to ensure the JVM has good performance. Detailed metrics including JVM CPU usage, Max Heap, Current Heap, Used Heap and Live Threads can all be tracked over time.
- [“RTView Servers Displays”](#): This series of displays is for monitoring the health of the RTView servers monitoring your system. RTView Manager metrics include connected state, number of clients and other status information for Data Server, Historian and Display Server processes.
- [“‘Drilldowns’ Displays”](#) or [“Other”](#) displays: These displays (such as [“Alerts Table by Component”](#) and [“Alert Detail for Component”](#)) are typically accessed by drilling down from other displays (however, [“Alerts History Table”](#) is accessed directly from the navigation tree).

Tomcat Displays

The Tomcat HTML displays provide extensive visibility into the health and performance of Tomcat application servers and installed web modules. The following Tomcat Views (and their associated displays) can be found under **Components** tab > **Application/Web Servers** > **Tomcat**.

Tomcat has the following displays:

- [“Tomcat Overview”](#)
- [“Tomcat Servers Heatmap”](#): Performance metrics for one Tomcat Server, including current and historic performance metrics.
- [“Single Tomcat Server”](#): Heatmap of performance metrics for all Web modules for one Tomcat Server.
- [“Tomcat Applications Heatmap”](#): Table and trend graphs of performance metrics for Web modules.
- [“Single Tomcat App”](#): Table and trend graphs of performance metrics for a single Web module.

Tomcat Overview

The Tomcat Overview is the top-level display for the Tomcat Solution Package, which provides a good starting point for immediately getting the status of all your Tomcat servers, web modules and connections. You can select the RTView DataServer for which you want to see data and easily view the current data for that DataServer including:


- The total number of active alerts for the selected DataServer, including the total number of critical and warning alerts.
- The greatest number of active sessions, top accesses per second, highest number of connections and the top process rate.
- A visual list of the top 10 servers with the greatest values for **Accesses**, **Requests**, **Cache Hit Rate**, **Process Rate**, **Sent** and **Received Rate**.

You can hover over each region in the upper half of the Overview to see more detail in a Summary display.

For example, clicking on the alerts in the CRITICAL and WARNING alerts region opens the **Alerts Table by Components** display.

The bottom half of the display provides a trend graph which traces **Active Sessions**, **Requests** per second and **Processing Time**.

You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.

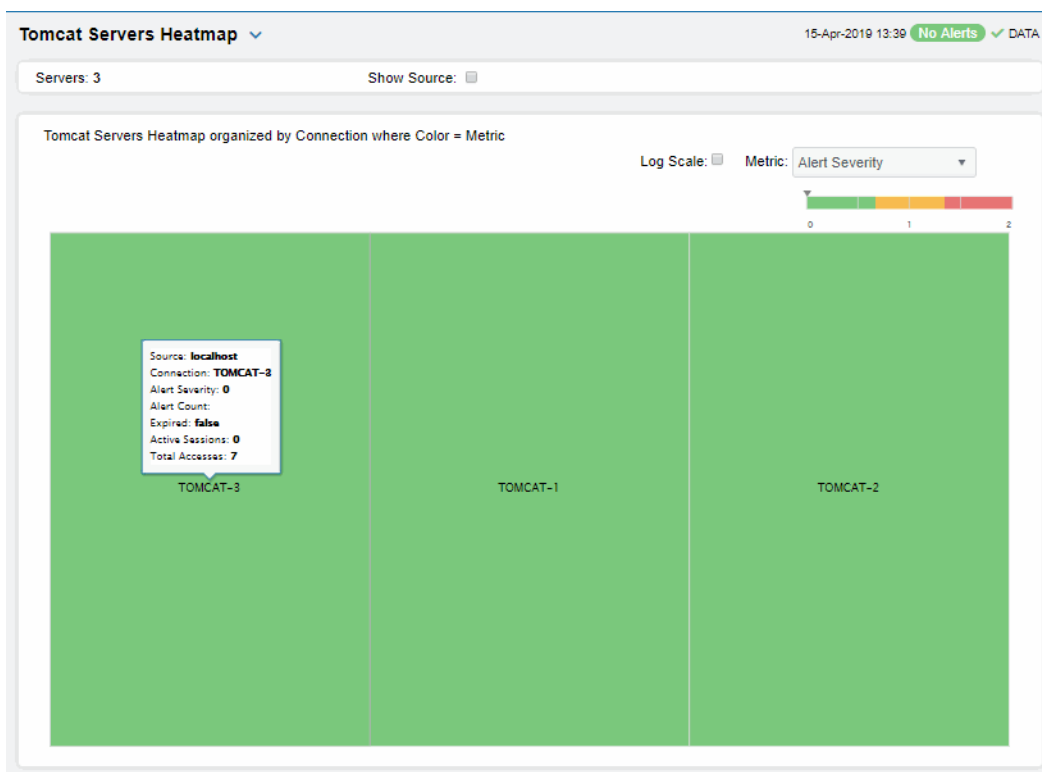


Tomcat Servers Heatmap

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

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

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

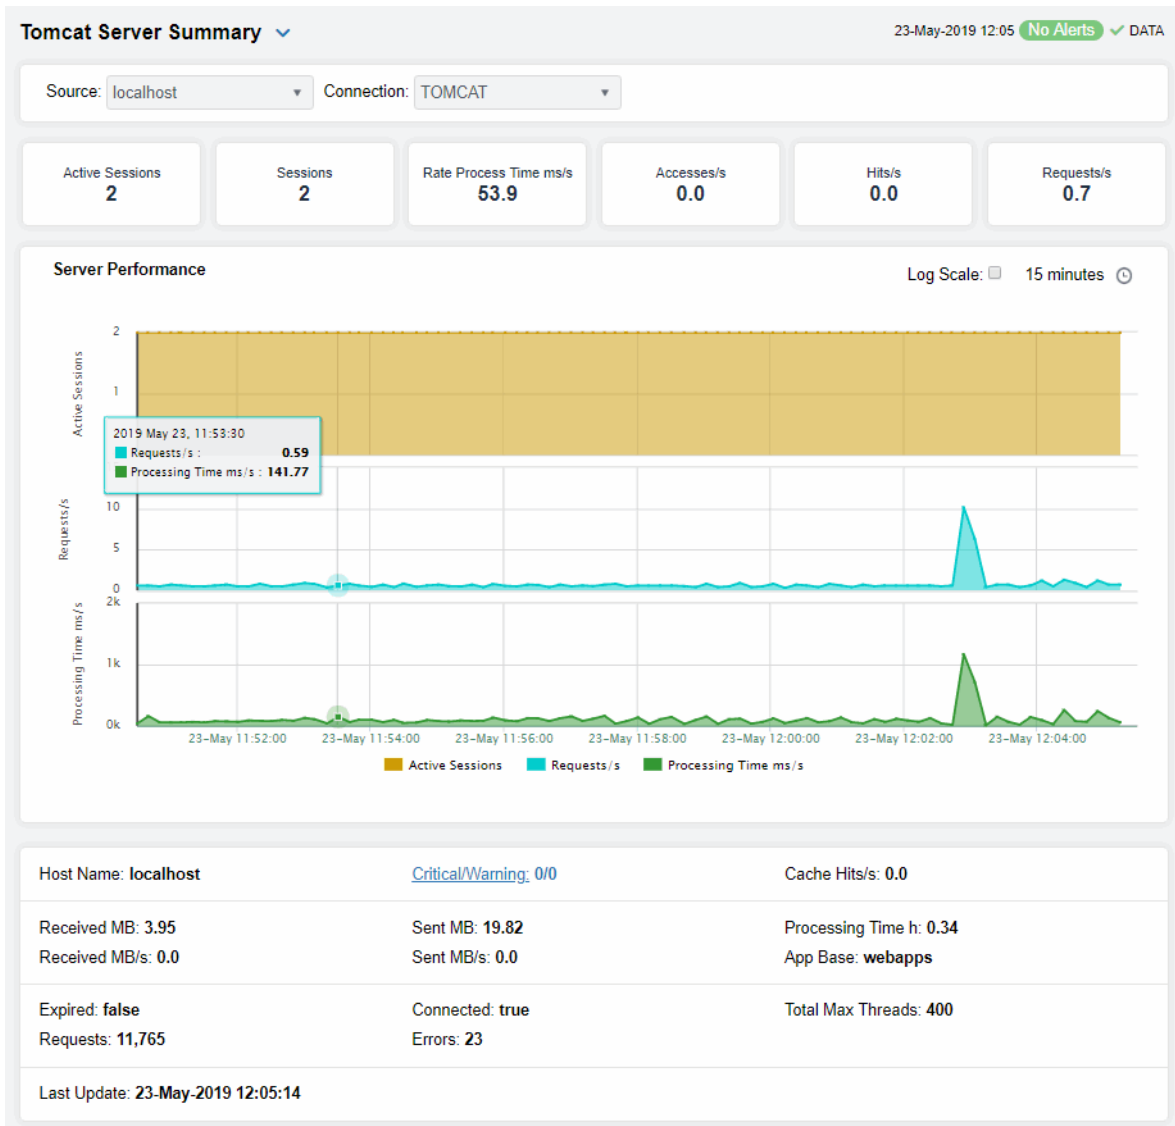


Single Tomcat Server

Track utilization and performance metrics for a connection on a Tomcat server. Clicking on the sessions/processing rate information boxes at the top of the display takes you to the **Tomcat Servers Table** display, where you can compare and sort performance values against other Tomcat servers.

The trend graph traces for **Processing Time per second**, **Requests per second** and (number of) **Active Sessions**. You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.

Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display.



Tomcat Applications

Investigate detailed utilization metrics for all Tomcat applications. This display contains all metrics available for Tomcat applications, including the total **Alert Count**, **Accesses/per second** and **Total Sessions**.

Choose a particular **Source** or **All**, and a particular **Connection** or **All**, from the drop-downs. Each row in the table contains data for a particular web module. You can search, filter, sort and choose columns to include by clicking a column header icon (to the right of each column label) and selecting **Filter**, **Sort Ascending**, **Sort Descending** or **Columns**.

Or just click a column header to sort.

Right-click on a table cell to **Export to Excel** or **Copy Cell Value**.

Double-click a web module to see details in the **Tomcat Application Summary display**.

Tomcat Applications

16-May-2019 15:39 No Alerts DATA

Source: localhost Connection: RTVIEW-SERVER-X

Applications: 6

Web Module	Expired	Alert Level	Alert Count	Total Accesses	Accesses/s	Total Sessions	Active Sessions	Expired Sessions	Cache Hits
/manager		✓		7	0.0	0	0	0	
/rtvdisplayplus		✓		7	0.0	0	0	0	
/host-manager		✓		7	0.0	0	0	0	
/docs		✓		7	0.0	0	0	0	
/rtvdataplus		✓		7	0.0	0	0	0	
/rtvquery		✓		7	0.0	0	0	0	

Tomcat Applications Heatmap

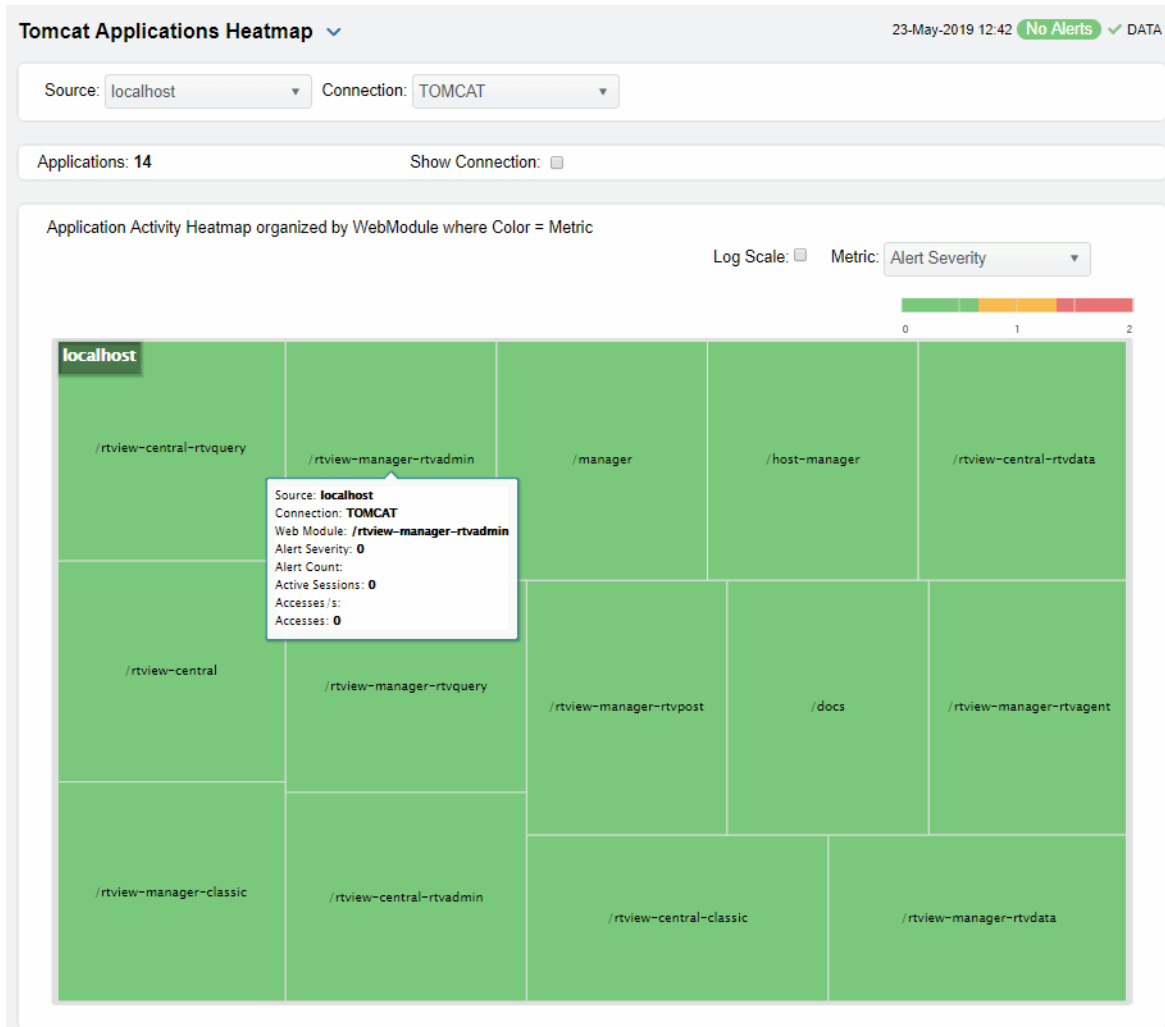
This heatmap allows you to view the status and alerts of Tomcat applications on a particular host or **All** hosts, and a particular connection or **All** connections.

Use the **Metric** drop-down menu to view the **Alert Severity**, **Alert Count**, **Active Sessions**, **Accesses per Second** or (the total number of) **Accesses**.

Each rectangle in the heatmap represents a web module. The rectangle color indicates the most critical alert state. Click on a rectangle to drill-down to the **Tomcat Application Summary** display and view metrics for a particular web module. Toggle between the commonly accessed Table and Heatmap displays by clicking the drop down list on the display title.

Mouse-over rectangles to view more details about host performance and status.

You can view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.



Single Tomcat App

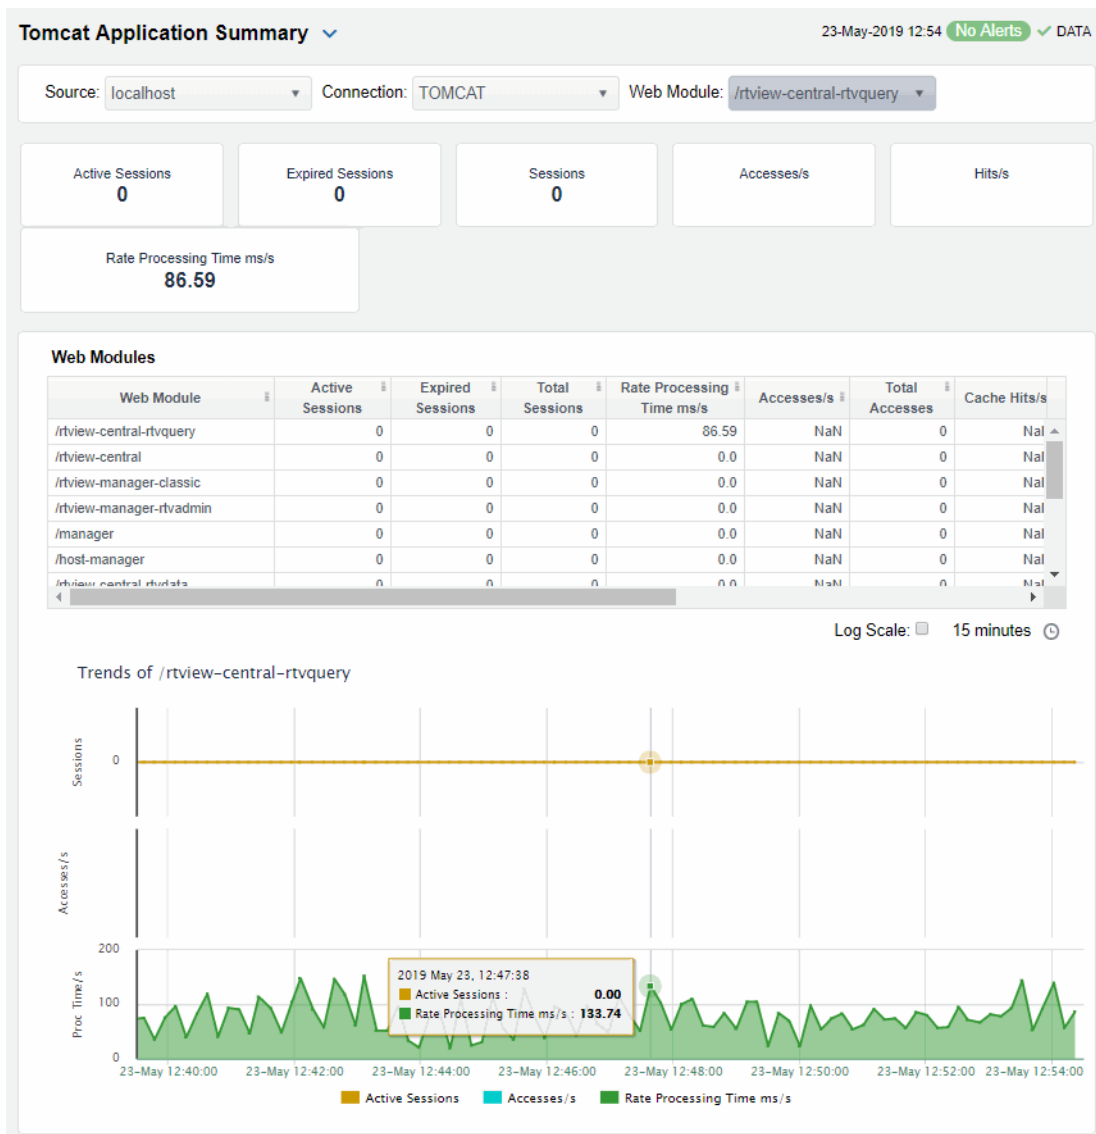
Track utilization and performance metrics for a particular Tomcat web module. Clicking on the sessions/processing rate information boxes at the top of the display takes you to the **Tomcat Servers Table** display, where you can compare and sort performance values against other Tomcat servers.

Use the **Web Modules** table to compare detailed utilization metrics for all web modules. Each row in the table contains data for a particular web module. You can search, filter, sort and choose columns to include by clicking a column header icon (to the right of each column label) and selecting **Filter**, **Sort Ascending**, **Sort Descending** or **Columns**.

Or just click a column header to sort. Right-click on a table cell to **Export to Excel** or **Copy Cell Value**.

The trend graph traces for **Processing Time per second**, **Accesses per second** and (the number of) **Active Sessions**. You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.

Clicking the **Critical/Warning** link at the bottom of the display opens the **Alerts Table by Component** display.



JVM Processes Displays

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

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

The HTML version features an overview display, ["JVM Overview"](#) (pictured below), and the following displays which can be found under **Components** tab > **Processes /JVM Processes** once RTView Manager is installed:

- ["JVMs Heatmap"](#): Heatmap of alert states for all JVM connections
- ["JVM Summary"](#): Table of connection details for all JVM connections.
- ["JVM System Properties"](#): Table of connection details for a single JVM as well as performance trend graphs.
- ["JVM GC Trends"](#): Trend graphs of garbage collection memory utilization.

JVM Overview

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

Choose a DataServer for which you want to see data and easily view the current data for that DataServer including:

- The total number of active alerts, including the total number of critical and warning alerts.
- The number of JVMs and the **Top CPU %** user across all servers.
- The maximum memory used and maximum number of threads.

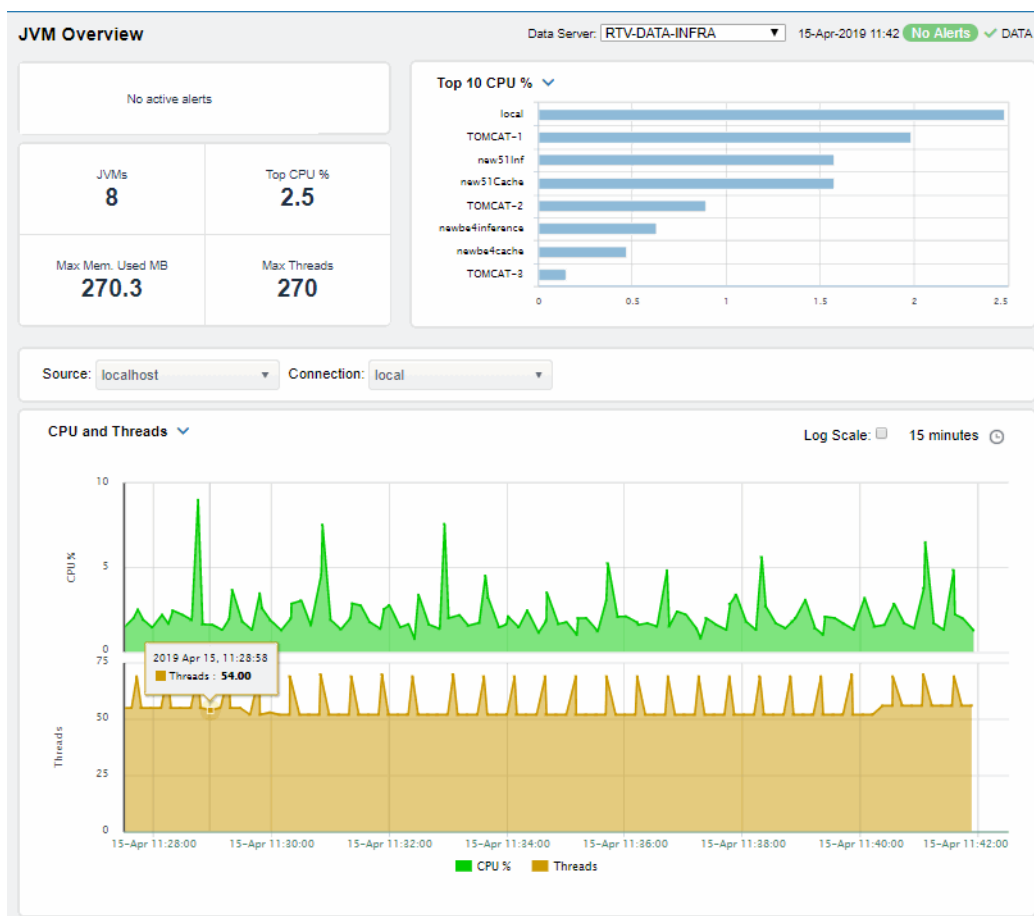
A bar graph shows the JVMs with **Top 10 CPU %** utilization. Use the drop-down menu to show JVMs with **Top 10 Used Heap** memory utilization and JVMs with **Top 10 Live Threads**.

You can hover over each region in the upper half of the Overview to see more detail. You can also drill down to see even more detail by clicking on each respective region in the Overview.

For example, clicking on the alerts in the CRITICAL and WARNING alerts region opens the Alerts Table display. Clicking on **Top CPU %** opens the ["JVM Summary"](#) display.

The bottom half of the display provides a performance trend graph for a connection on the DataServer. Choose a **Source** and **Connection** from the drop-down menus. Use the trend graph drop-down menu to show metrics for **CPU and Threads** utilization or Heap Memory utilization.

You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.



JVMs Table

Investigate JVM connection utilization metrics and configuration information for one or all JVMs. Choose one or **All** JVMs from the **Source** drop-down menu. Each row in the table contains data for a particular connection on the selected JVM(s).

This display contains all metrics available for JVM connections, including the **Port** number and the current most critical **Alert Level**, where:

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

You can search, filter, sort and choose columns to include by clicking a column header icon (to the right of each column label) and selecting **Filter**, **Sort Ascending**, **Sort Descending** or **Columns**.

Or just click a column header to sort.

Right-click on a table cell to **Export to Excel** or **Copy Cell Value**.

Double-click on a table row to drill-down to the **JVM Summary - HTML** display and view metrics for the JVM hosting the connection.

Check the **Show Inactive** box to include inactive connections.

Toggle between the commonly accessed **Table** and **Heatmap** displays by clicking the drop down list on the display title.

Right-click on a table cell to **Export to Excel** or **Copy Cell Value**.

JVMs Table 15-Apr-2019 12:59 No Alerts DATA

Source: - All -

Active Only: ☐ JVMs: 14

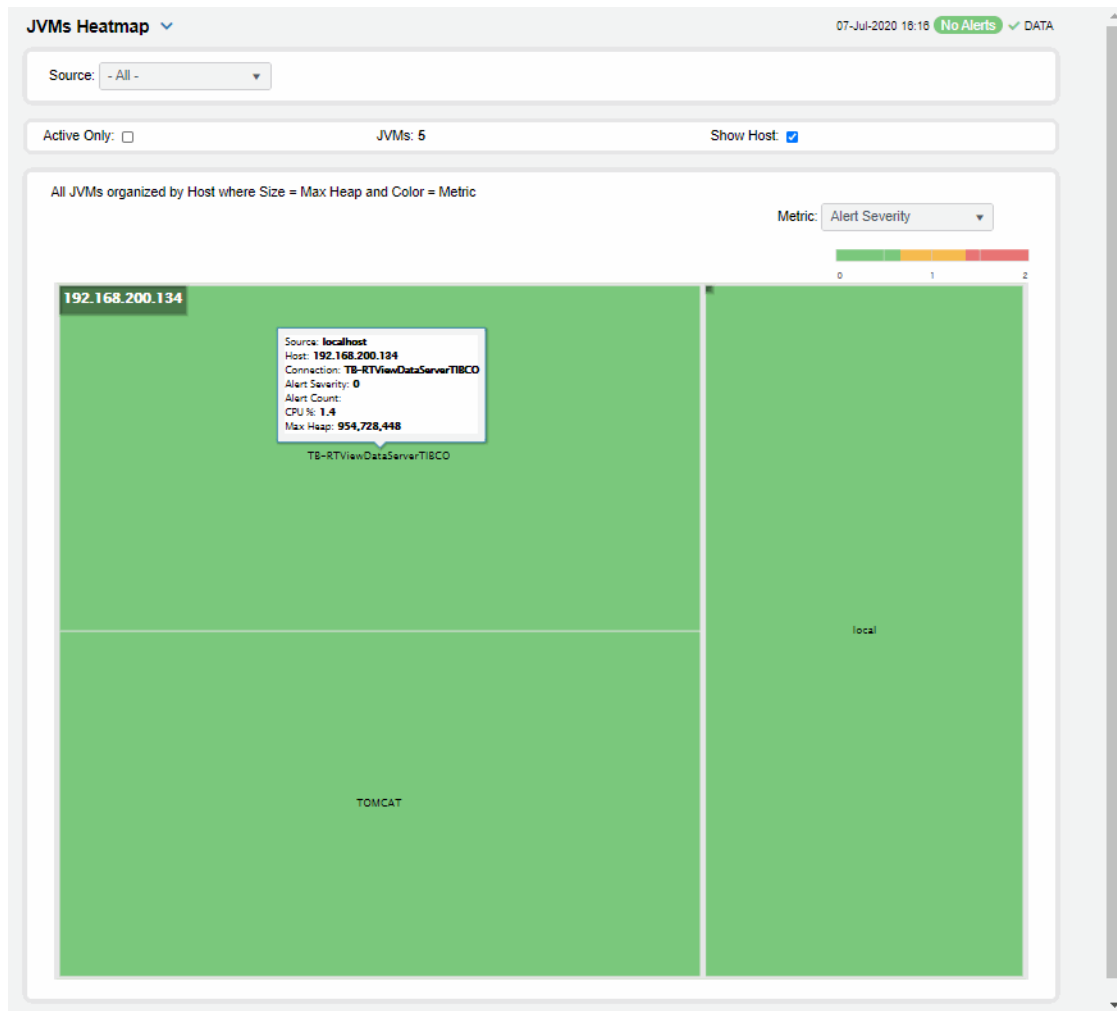
Source	Connection	Expired	Connected	Alert Level	Alert Count	Host	Port	PID	CPU %
localhost	TOMCAT-3		✓	✓		172.30.1.85	9999		0.
localhost	TOMCAT-2		✓	✓		172.30.1.57	9999		1.
localhost	TOMCAT-1		✓	✓		172.30.1.174	9999		1.
localhost	local		✓	✓				348@172.30.1.104	1.
localhost	local		✓	✓				9@172.30.1.174	1.
localhost	local		✓	✓				1114@172.30.1.31	1.
localhost	local		✓	✓				1120@172.30.1.31	1.
localhost	local		✓	✓				123@172.30.1.97	1.
localhost	local		✓	✓				126@172.30.1.97	1.
localhost	local		✓	✓				9@172.30.1.174	1.
localhost	new51Inf		✓	✓		192.168.200.144	58701		1.
localhost	newbe4cache		✓	✓		192.168.200.144	48700		0.
localhost	new51Cache		✓	✓		192.168.200.144	58700		0.
localhost	newbe4inference		✓	✓		192.168.200.144	48701		0.

JVMs Heatmap

View the alert state for all monitored JVM connections for one or all sources, as well as CPU and memory utilization. The heatmap organizes JVM connections by host, and uses color to show the selected **Metric** value for each JVM connection.

Each rectangle in the heatmap represents a JVM connection. The rectangle size represents the amount of memory reserved for that process; a larger size is a larger value. Each Metric (selected from the drop-down menu) has a color gradient bar that maps relative values to colors.

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



Metric

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

Alert Severity

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





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

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

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

Alert Count

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

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.

JVM Summary

Track utilization by a single connection on a JVM, including **Memory** and **CPU** usage, amount of **Committed Memory** (the amount of memory, in megabytes, guaranteed to be available for use by the JVM).

Choose a **Source** and **Connection** from the drop down menus. The amount of committed memory can be a fixed or variable size. If set to be a variable size, the amount of committed memory can change over time, as the JVM may release memory to the system. This means that the amount allocated for **Committed** memory could be less than the amount initially allocated. Committed memory will always be greater than or equal to the amount allocated for **Used memory** and **Maximum Memory** used, number of **Threads** and **Peak Threads**.

You can also verify whether the memory usage has reached a plateau. And if usage is approaching the limit, determine whether to allocate more memory.


Clicking on the information boxes at the top of the display takes you to the **JVMs Table - HTML** display, where you can view and compare with other connections on the JVM.

You can set the time range for the trend graph to trace. You can also choose what to trace from the drop-down menu:

- **CPU and Threads** traces the amount of CPU used by the JVM and the total number of live threads.
- **Heap Memory** traces the amount of memory used for memory management by the application in the time range specified. This value may change or be undefined. Note that a memory allocation can fail if the JVM attempts to set the Used memory allocation to a value greater than the **Committed** memory allocation, even if the amount for **Used** memory is less than or equal to the Maximum memory allocation (for example, when the system is low on virtual memory).

At the bottom of the display you also can get JVM operating system information, the number of processors available to the JVM, the **Architecture** which is the ISA used by the processor, the number of **Daemon Threads** and **Input Arguments** for starting JVM.

Clicking the **Critical/Warning** link at the bottom of the display opens the Alerts Table by Component display.



JVM System Properties

View JVM arguments in the RuntimeMXBean InputArguments attribute, command line arguments for starting applications and system properties settings for a connection.

Choose a **Source** and **Connection** from the drop-down menus. You can search, filter, sort and choose columns to include by clicking a column header icon (to the right of each column label) and selecting **Filter**, **Sort Ascending**, **Sort Descending** or **Columns**.

Or just click a column header to sort.

Right-click on a table cell to **Export to Excel** or **Copy Cell Value**.

Click a column header to sort column data in ascending or descending order or right-click to filter data shown in the display. Toggle between the commonly accessed Table and Heatmap displays by clicking the drop down list on the display title.

JVM System Properties 15-Apr-2019 13:05 No Alerts DATA

Source: localhost Connection: local

JVM Arguments

Value
-Dcom.sls.rtv.customRtvAppManagerClassName=com.sls.jmsrtvutils.RtvAppManager
-Dcom.sun.management.jmxremote.authenticate=false
-Dcom.sun.management.jmxremote.ssl=false
-Dcom.sls.rtv.useLog4j=true
-Dcom.sls.rtv.RTVLog4jLevel=info
-Dcom.sls.rtv.showLogCategory=true
-Xmx1024m
-Xms256m

Command-Line Arguments

Value

System Properties

Property	Value
com.sun.management.jmxremote.rmi.port	3288
awt.toolkit	sun.awt.windows.WToolkit
com.sls.rtv.RTVLog4jLevel	info
file.encoding.pkg	sun.io
java.specification.version	1.7
sun.cpu.isalist	amd64

JVM GC Trends

Track JVM garbage collection memory utilization trends for a single connection. Choose a **Source**, **Connection** and **Garbage Collector** from the drop-down menus. The upper trend graph traces the following for the selected garbage collector for the time range specified:

- **Max:** The maximum amount of memory, in megabytes, used for JVM garbage collection.
- **Committed:** The amount of memory, in megabytes, guaranteed to be available for use by JVM non-heap memory management. Note that the amount of committed memory can be a fixed or variable size. If set to be a variable size, it can change over time, as the JVM may release memory to the system. This means that the amount allocated for committed memory could be less than the amount initially allocated. Committed memory will always be greater than or equal to the amount allocated for used memory.
- **Used - Before:** The amount of memory, in megabytes, used before the last garbage collection.
- **Used - After:** The amount of memory, in megabytes, used after the last garbage collection.

The lower trend graph traces the following for the selected garbage collector for the time range specified:

- **Duration:** The duration, in seconds, of garbage collection.
- **Duty Cycle:** The percentage of time that the application spends in garbage collection.

You can hover over the trend graph to see the values at a particular time.

You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.



RTView Servers Displays

The following RTView Servers displays can be found under **Components** tab > **Processes** > **JVM Processes** > **RTView Servers** after installation.

These displays present performance data for all RTView servers. Use these displays to monitor the health of the servers monitoring your system. Displays are:

- **"Data Servers"**: Shows metrics for RTView Data Servers.
- **"Data Server Summary"**: Shows details for one Data Server.
- **"Display Servers"**: Note that this display does not contain data.
- **"Display Server Summary"**: Note that this display does not contain data.
- **"Historian Servers"**: Shows metrics for RTView Historian Servers.

Data Servers

View connections on one or all RTView Data Servers, as well as connection status and client count. Choose one or **All** data servers from the **Source**: drop-down menu. Each row in the table is a different data server.

You can search, filter, sort and choose columns to include by clicking a column header icon (to the right of each column label) and selecting **Filter**, **Sort Ascending**, **Sort Descending** or **Columns**.

Or just click a column header to sort.

Right-click on a table cell to **Export to Excel** or **Copy Cell Value**.


Double-click to drill-down to details about the selected data server as well as the selected connection in the **Data Server Summary** display.

Source	Connection	Client Count	Serving Data	Expired	time_stamp
localhost	local	7	✓		16-Apr-2019 14:56:16

Data Server Summary

Track utilization metrics for a specific data server and a connection. Choose a data server and a connection from the **Source** and **Connection** drop-down menus. View client details such as client ID, IP address, process name, host and duration. The trend graph traces data sent and client count for the selected connection. You can hover over the trend graph to see the values at a particular time. You can specify the

time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.



Historian Servers

Track the status of RTView Historian Servers, their connections, status and role and data configuration file usage. View the caches that are archived by the Historian application, substitution variables associated with the history cache configuration file, as well as the history cache status.

Each row in the table contains data for a particular server. You can search, filter, sort and choose columns to include by clicking a column header icon (to the right of each column label) and selecting **Filter**, **Sort Ascending**, **Sort Descending** or **Columns**. Or just click a column header to sort. Right-click on a table cell to **Export to Excel** or **Copy Cell Value**.

RTView Historians Table 16-May-2019 15:12 ✓ DATA

Source: localhost

Historians: 2

Source	Connection	Connected To DB	Suspended	Config File Count	Primary	Expired	Time Stamp
localhost	RTVMGR-HIST-X2			0	✓		16-May-2019 15:12:26
localhost	EMSMON-HIST-3			0	✓		16-May-2019 15:12:26

‘Drilldowns’ Displays

This View contains the following displays:

- [“Alerts History Table”](#): Track history of any alert that has occurred in your RTView system.
- [“Alerts Table by Component”](#): Track alerts associated with CIs shown in a display.
- [“Alert Detail for Component”](#): Investigate an alert instance and its history.

Alerts History Table

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

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

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

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

You can search, filter, sort and choose columns to include by clicking a column header icon (to the right of each column label) and selecting **Filter**, **Sort Ascending**, **Sort Descending** or **Columns**. Right-click on a table cell to **Export to Excel**.

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

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

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

Alert Level	Ack	Cleared	Alert Name	Alert Index	Alert Text	Owner	Id	Source	Row Update Time
			JvmMemoryUsed-High	win44-newbe4cache	High Alert Limit exceeded		144843		2019-Oct-07 09:11:41
		<input checked="" type="checkbox"/>	JvmMemoryUsed-High	win44-newbe4cache	High Alert Limit exceeded		144843		2019-Oct-07 09:11:50
			BwServerCpuUsed-High	sl4-64(simon)	High Alert Limit exceeded		144911		2019-Oct-07 09:18:4
		<input checked="" type="checkbox"/>	BwServerCpuUsed-High	sl4-64(simon)	High Alert Limit exceeded		144911		2019-Oct-07 09:19:0
			BwProcessAvgElapsedTime-High	sl4-64(simon)-domains	High Alert Limit exceeded		144855		2019-Oct-07 09:10:5
			BwProcessAvgElapsedTime-High	sl4-64(simon)-domains	High Alert Limit exceeded		144854		2019-Oct-07 09:10:5
			BwProcessAvgElapsedTime-High	sl4-64(simon)-domains	High Alert Limit exceeded		144853		2019-Oct-07 09:10:5
			BwProcessAvgElapsedTime-High	sl4-64(simon)-domains	High Alert Limit exceeded		144852		2019-Oct-07 09:10:5
			BwProcessAvgElapsedTime-High	sl4-64(simon)-domains	High Alert Limit exceeded		144851		2019-Oct-07 09:10:5
			BwProcessAvgElapsedTime-High	sl4-64(simon)-domains	High Alert Limit exceeded		144856		2019-Oct-07 09:11:5
			BwProcessAvgElapsedTime-High	slhpux11(simon)-domain	High Alert Limit exceeded		144865		2019-Oct-07 09:12:1
			BwProcessAvgElapsedTime-High	slhpux11(simon)-domain	High Alert Limit exceeded		144864		2019-Oct-07 09:12:1
			BwProcessAvgElapsedTime-High	slhpux11(simon)-domain	High Alert Limit exceeded		144863		2019-Oct-07 09:12:1
			BwProcessAvgElapsedTime-High	slhpux11(simon)-domain	High Alert Limit exceeded		144862		2019-Oct-07 09:12:1
			BwProcessAvgElapsedTime-High	slhpux11(simon)-domain	High Alert Limit exceeded		144861		2019-Oct-07 09:12:1
			BwProcessAvgElapsedTime-High	slhpux11(simon)-domain	High Alert Limit exceeded		144860		2019-Oct-07 09:12:1
			BwProcessAvgElapsedTime-High	slhpux11(simon)-domain	High Alert Limit exceeded		144859		2019-Oct-07 09:12:1
			BwProcessAvgElapsedTime-High	sl4-64(simon)-domains	High Alert Limit exceeded		144872		2019-Oct-07 09:12:1
			BwProcessAvgElapsedTime-High	sl4-64(simon)-domains	High Alert Limit exceeded		144871		2019-Oct-07 09:12:1
			BwProcessAvgElapsedTime-High	sl4-64(simon)-domains	High Alert Limit exceeded		144870		2019-Oct-07 09:12:1
			BwProcessAvgElapsedTime-High	sl4-64(simon)-domains	High Alert Limit exceeded		144869		2019-Oct-07 09:12:1
			BwProcessAvgElapsedTime-High	sl4-64(simon)-domains	High Alert Limit exceeded		144868		2019-Oct-07 09:12:1

Alerts Table by Component

As an alternative to the **Alerts Table**, use the **Alerts Table by Component** to track and manage all alerts that are specifically associated with the CIs shown in a display.

You access the **Alerts Table by Component** by clicking (the alert status icon) in the title bar of other displays. The display in which you click is the source display.

Package provides the technology label associated with the alerts shown. For example, **Jvm**, **Tomcat** and **Host** are the technology labels for Java Virtual Machines, Tomcat applications and servers (respectively). These labels are also correlated with the RTView solution package names (for example, the Solution Package for Host Agent). **Category** lists all alert categories related to the source display.

Use the **ACK** and **Cleared** drop-downs to filter the table by **All**, **True** or **False**.

See the **Alert Level** column icon, where:



The alert reached its ALARM LEVEL threshold in the table row.



The alert reached its WARNING LEVEL threshold in the table row.

To investigate, click:

Alert Detail

to open the **Alert Detail for Component** where you can see the current and historical conditions that precipitated the alert being executed.

[Go to CI](#) to open the summary display for the CI associated with the alert where you can investigate utilization metrics for the CI leading up to the alert being executed.

You can search, filter, sort and choose columns to include by clicking a column header icon (to the right of each column label) and selecting **Filter**, **Sort Ascending**, **Sort Descending** or **Columns**. Right-click on a table cell to **Export to Excel**. Use **Ctrl** + click or **Shift** + click to select multiple alerts.

With one or more alerts selected, click [Own](#) to set the alert(s) owner field, [Acknowledge](#) to acknowledge the alert(s), [Unacknowledge](#) to clear the acknowledgement on previously acknowledged alert(s), [Add Comment](#) to add a comment to the alert(s).

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

Alerts Table by Component 02-May-2019 11:05:09 DATA OK

Package: Host Category: CPU:Network:Storage Cleared: False ACK: False

Alert Count: 16

Row	Update Time	Acknowledge	Cleared	Alert Level	Alert Name	Alert Index Values	
2018-Nov-09 23:54:0				🔔	HostCpuPercentHigh	SL-DEMO;SLHOST16(sl_qa)	High V
2018-Oct-01 06:20:10				⚠️	HostCpuPercentHigh	SL-DEMO;SLHOST17(sl_amx)	High A
2019-May-02 03:28:5				🔔	HostMemoryUsedHigh	SL-DEMO-LX;192.168.200.92	High V
2018-Oct-01 06:19:38				⚠️	HostVirtualMemoryUsedH	SL-DEMO;SLHOST17(sl_amx)	High A
2018-Oct-01 06:18:38				🔔	HostMemoryUsedHigh	SL-DEMO;SLHOST17(sl_amx)	High V
2018-Jan-12 11:38:56				⚠️	HostCpuPercentHigh	SL-DEMO-LX;192.168.200.205	High A
2019-May-02 10:40:3				⚠️	HostVirtualMemoryUsedH	SL-DEMO-LX;192.168.200.42	High A
2019-Apr-25 10:19:43				🔔	HostMemoryUsedHigh	SL-DEMO;SLHOST8	High V
2018-Jun-19 09:22:23				⚠️	HostCpuPercentHigh	SL-DEMO-LX;192.168.200.202	High A
2018-Nov-09 10:33:54				⚠️	HostVirtualMemoryUsedH	SL-DEMO;SLHOST16(sl_qa)	High A
2018-May-01 22:45:4				⚠️	HostCpuPercentHigh	SL-DEMO-LX;192.168.200.92	High A

[Alert Detail](#)
[Go to CI](#)
[Own](#)
[Acknowledge](#)
[Unacknowledge](#)

[Add Comment](#)
[Clear All Comments](#)

Alert Detail for Component

Use the **Alert Detail for Component** display to investigate current and historical activity of a specific alert instance as it applies to the associated CI, and also compare against **Metric History** trends of the associated CI. A trend graph for the CI associated with the alert instance. You can hover over the trend graph to see the values at a particular time. You can specify the time range for the trend graph and view data based on a log scale, which enables visualization on a logarithmic scale and should be used when the range in your data is very broad.

Access the **Alert Detail for Component** display by clicking  in the **Alerts Table** or  in the **Alerts Table by Component** display.

The **Alert History** table at the bottom of the display contains a row of data for each time the alert instance was updated. See the alert **ID**, **Row Update Time**, **Cleared** status and **Reason**, **Owner** and the **Alert Level** column icon, where:




The alert reached its ALARM LEVEL threshold in the table row.




The alert reached its WARNING LEVEL threshold in the table row.

You can search, filter, sort and choose columns to include by clicking a column header icon (to the right of each column label) and selecting **Filter**, **Sort Ascending**, **Sort Descending** or **Columns**. Right-click on a table cell to **Export to Excel**. Use **Ctrl** + click or **Shift** + click to select multiple alerts.

To investigate, click:

 to see utilization conditions for the CI associated with the alert in a summary display.

[Admin](#) to open the **Alert Configuration for Component** display where you can see, modify and refine alert threshold settings for that particular alert. A trend graph traces the relevant alert metric for the CI so you can adjust thresholds in real-time.



Alerts Displays

This section describes the [“Alerts Table”](#) for RTView Manager. For details about alerts for RTView Manager, see [“Alerts for RTView Manager”](#).

Alerts Table

Use this display to track and manage all alerts that have occurred in the system, where:



One or more alerts exceeded their ALARM LEVEL threshold in the table row



One or more alerts exceeded their WARNING LEVEL threshold in the table row

You can search, filter, sort and choose columns to include by clicking a column header icon (located to the right of each column label) and selecting **Filter**, **Sort Ascending**, **Sort Descending** or **Columns**. Use the **Ack'd** and **Cleared** drop-downs to filter the table by those columns. Right-click on a table cell to **Export to Excel** or **Copy Cell Value**. Use **Ctrl** + click or **Shift** + arrow to select multiple alerts. To investigate, select one alert and click:

Details

to open the **Component Alert Detail** display to get details about that particular alert instance as it specifically applies to the associated CI.

CI

to see utilization conditions for the CI associated with the alert during the seconds (minutes, hours or days) leading up to the alert being executed in a summary display.

With one or more alerts selected, you can click **Own** to set the alert(s) owner field, **Ack** to acknowledge the alert(s), **Unack** to clear the acknowledgement on previously acknowledged alert(s), **Clear** to set the **Cleared** flag on the selected alert(s), **Comment** to add a comment to the alert(s) and **CI** to get details about the CI associated with the alert (these buttons are enabled when you click one or more alerts).

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

Alerts Table											
30-Apr-2019 13:47:46 ✓ DATA											
<div> <div>Own Ack Unack Clear Comment Details CI</div> <div>Ack'd: all Cleared: false Cmdb Filter: ***** Alert Count: 92</div> </div>											
Time	Ack	Clr	Sevl	Alert Name	Alert Text	Owner	ID	Source	Comments	CI	
2019-Apr-30 00:04:07			⚠	JvmNotConnected	Server disconnected		1043	RTV-DATA-TIB		win4	
2019-Apr-30 01:34:49			⚠	JvmNotConnected	Server disconnected		1009	Z-SIMDATA-1		local	
2019-Apr-30 01:34:49			⚠	JvmNotConnected	Server disconnected		1008	Z-SIMDATA-1		local	
2019-Apr-30 01:34:49			⚠	JvmNotConnected	Server disconnected		1007	Z-SIMDATA-1		local	
2019-Apr-30 01:34:49			⚠	JvmNotConnected	Server disconnected		1006	Z-SIMDATA-1		local	
2019-Apr-30 01:34:49			⚠	JvmNotConnected	Server disconnected		1005	Z-SIMDATA-1		local	
2019-Apr-30 01:34:49			⚠	JvmNotConnected	Server disconnected		1004	Z-SIMDATA-1		local	
2019-Apr-30 01:34:49			⚠	JvmNotConnected	Server disconnected		1003	Z-SIMDATA-1		local	
2019-Apr-30 01:34:49			⚠	JvmNotConnected	Server disconnected		1002	Z-SIMDATA-1		local	
2019-Apr-30 01:34:49			⚠	JvmNotConnected	Server disconnected		1001	Z-SIMDATA-1		local	
2019-Apr-30 01:34:49			⚠	JvmNotConnected	Server disconnected		1000	Z-SIMDATA-1		local	
2019-Apr-30 12:01:02			⚠	JvmCpuPercentHigh	High Alert Limit exceed		1064	Z-SIMDATA-1		local	
2019-Apr-30 13:44:01			🔔	JvmCpuPercentHigh	High Warning Limit exc		928739	RTV-DATA-KAF		Insta	
2019-Apr-30 13:47:04			🔔	JvmCpuPercentHigh	High Warning Limit exc		928747	RTV-DATA-KAF		Insta	
2019-Apr-30 01:36:49			🔔	HostCpuPercentHigh	High Warning Limit exc		1010	Z-SIMDATA-1		defa	
2019-Apr-30 01:36:49			🔔	HostCpuPercentHigh	High Warning Limit exc		1010	Z-SIMDATA-1		defa	
2019-Apr-30 02:05:10			⚠	HostCpuPercentHigh	High Alert Limit exceed		1011	Z-SIMDATA-1		defa	

For details about alerts for RTView Manager, see ["Alerts for RTView Manager"](#).

Admin Displays

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

- **"Alert Administration"**: Displays active alerts and provides interface to modify, enable and manage alerts.
- **"Alert Overrides Admin"**: Set and modify alert overrides. Access this display from the **Alert Administration** display.
- **"Cache Table"**: View cached data that RTView is capturing and maintaining, and use this data use this for debugging with SL Technical Support.

For details about alerts for RTView Manager, see ["Alerts for RTView Manager"](#).

Alert Administration

The **Alert Administration** display allows administrators to enable/disable alerts and manage alert thresholds. The table describes the global settings for all alerts on the system.

You can set the **Delay** time (the number of seconds that must pass before an alert is triggered, where **0** sets it to immediately execute).

You can set the **Warning Level** which executes a single warning alert when the number of seconds specified here is exceeded. To set the warning to occur sooner, reduce the **Warning Level** value. To set the warning to occur later, increase the **Warning Level** value.

You can set the **Alarm Level** which executes a single alarm alert when the number of seconds specified here is exceeded. To set the alarm to occur sooner, reduce the **Alarm Level** value. To set the alarm to occur later, increase the **Alarm Level** value.

Note: For low value-based alerts (an alert that executes based on a value going below a certain threshold), to set the alarm to occur sooner you increase the **Alarm Level** value. To set the alarm to occur later, reduce the **Alarm Level** value.

You can apply alert thresholds globally or as an *override*. Setting override alerts allows you to set thresholds for a subset of your resources, or for a single resource (for example, a single server). Override alerts are useful if the majority of your resources require the same threshold setting, but there are a few resources that require a different threshold setting. For example, you might not usually be concerned with execution time at a process level, but perhaps certain processes are critical. In this case, you can apply alert thresholds to each process individually. See below for instructions.

You can filter, sort and choose columns to include by clicking a column header icon (located to the right of each column label) and selecting **Filter**, **Sort Ascending**, **Sort Descending** or **Columns**. Use the **Ack'd** and **Cleared** drop-downs to filter the table by those columns. Right-click on a table cell to **Export to Excel**.

To set thresholds and enable a global alert:

Select an alert and, under **Settings for alert** (in the lower portion of the screen), modify settings for the alert **Delay**, **Warning Level** and/or **Alarm Level** and **Save Settings**. With that alert selected, check the **Alert Enabled** box under **Settings for alert** (in the lower portion of the screen) and **Save Settings**. The **Alert Enabled** box (next to the selected alert) is now checked.

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

To set thresholds and enable an override alert:

To set an override alert, select an alert and click **Override Settings** to open the **Alert Overrides Admin** display.

Alerts Administration 30-Apr-2019 10:34:01 ✓ DATA OK

Package: All http://rtvdemos.sl.com/emdemo_central_rtquery

Alert Name	Alert Enabled	Alert Delay	Warning Level	Alert Level	Override Count
HostNetworkTxRateHigh	<input type="checkbox"/>	30	50	75	0
HostProcessCountLow	<input type="checkbox"/>	30	15	5	0
HostStateData	<input type="checkbox"/>	30			0
HostStorageUsedHigh	<input type="checkbox"/>	30	80	90	0
HostSwapUsedHigh	<input type="checkbox"/>	30	75	90	0
HostVirtualMemoryUsedHigh	<input type="checkbox"/>	30	75	90	0
JvmCpuPercentHigh	<input checked="" type="checkbox"/>	60	50	70	0
JvmGcDutyCycleHigh	<input type="checkbox"/>	30	50	75	0
JvmMemoryUsedAfterGCHigh	<input type="checkbox"/>	0	1	80	0
JvmMemoryUsedHigh	<input checked="" type="checkbox"/>	60	75	86	0
JvmNotConnected	<input checked="" type="checkbox"/>	60			0
JvmStateData	<input type="checkbox"/>	30			0
JvmThreadCountHigh	<input checked="" type="checkbox"/>	60	8000	12000	0

Page 2 of 5 101 - 200 of 432 items

Settings for alert

Alert Enabled: ☐ Delay: Warning Level: Alert Level:

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

For details about alerts for RTView Manager, see ["Alerts for RTView Manager"](#).

Also see ["Alert Overrides Admin"](#).

Alert Name	The name of the alert.
Alert Enabled	When checked, the alert is enabled globally.
Alert Delay	The amount of time (in seconds) that the value must be above the specified Warning Level or Alarm Level threshold before an alert is executed. 0 is for immediate execution.

Warning Level	The global warning threshold for the selected alert. When the specified value is exceeded a warning is executed.
Alert Level	The global alarm threshold for the selected alert. When the specified value is exceeded an alarm is executed.
Override Count	The number of times thresholds for this alert have been defined individually in the Tabular Alert Administration display. A value of: -0 indicates that no overrides are applied to the alert. -1 indicates that the alert does not support overrides.

Settings for alert

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

Alert Enabled	Check / uncheck this box to enable or disable the selected alert globally.
Delay	Enter the amount of time (in seconds) that the value must be above the specified Warning Level or Alarm Level threshold before the selected alert is executed. 0 is for immediate execution.
Warning Level	Enter the global warning threshold for the selected alert. When the specified value is exceeded a warning is executed. To set the warning to occur sooner, reduce the Warning Level value. To set the warning to occur later, increase the Warning Level value.
Alert Level	Enter the global alarm threshold for the selected alert. When the specified value is exceeded an alarm is executed. To set the alarm to occur sooner, reduce the Alarm Level value. To set the warning to occur later, increase the Alarm Level value. NOTE: For low value-based alerts (such as EmsQueuesConsumerCountLow), to set the alarm to occur sooner, increase the Alarm Level value. To set the alarm to occur later, reduce the Alarm Level value.
Save Settings	Click to apply alert settings for the selected alert.
Original Defaults	Click to revert to original alert settings for the selected alert.
Override Settings	Click to set an alert override in the Alert Overrides Admin display on the selected alert.

Alert Overrides Admin

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

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

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

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

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

You can filter, sort and choose columns to include by clicking a column header icon (located to the right of each column label) and selecting **Filter**, **Sort Ascending**, **Sort Descending** or **Columns**. Use the **Display** drop-down to filter the table to show **All** resources, only resources with the **Overridden** alert applied or **Free** resources (to show only resources without the alert override applied). Right-click on a table cell to **Export to Excel** or **Copy Cell Value**.

To set overrides:

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

To alter overrides:

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

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

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

Alert Overrides Administration 10-Jun-2020 18:25:07 DATA OK

Alert: SolBrokerEgressByteRateHigh Override Type: PerMsgRouter Display: All

Search: * RegEx: ☐

Connection	Override Enabled	Alert Enabled	Alert Delay	Warning Level	Alert Level
solaceLoaner					
VMR_144-SP					
VMR_209					
VMR_142-US-EAST-					
VMR_114-US-EAST-					
VMR_112-SP					

Settings for selected index

Override Enabled: ☒ Alert Enabled: ☒ Alert Delay: Warning Level:

Alert Level:

Add Override Save Settings Remove Override

Cache Table

View the raw data that RTView is capturing and maintaining to investigate utilization and capacity metrics, as well as connection details, for caches on a data server.

Select a **Data Server** from the drop-down menu. The upper table contains a row of data for each cache on the selected data server. You can see the current number of **Rows** and **Columns** in each table and the amount of **Memory** used. You can also find out the cache **Table** type of which there are five:

- **current** tables show the most recently received values for each index.
- **current_condensed** tables are current tables with primary compaction configured.
- **history** tables show the historical values for each index.
- **history_condensed** tables are history tables with primary compaction configured.
- **history_combo** tables are history tables with primary compaction configured, and which is also configured to store rows of recent raw data followed by rows of older condensed data.

Select a cache to see connection utilization details for that cache in the lower table. The lower table shows the contents of the selected cache table. Available columns vary by cache. For example, a JVM cache table might provide **BootClassPath** and **InputArgument** columns, and a Tomcat cache might provide **RateAccess** and **cacheMaxSize** columns.

You can search, filter, sort and choose columns to include by clicking a column header icon (to the right of each column label) and selecting **Filter**, **Sort Ascending**, **Sort Descending** or **Columns**. Or just click a column header to sort.

Right-click on a table cell to **Export to Excel** or **Copy Cell Value**. Use **Ctrl + click** or **Shift + click** to select multiple alerts. Use **History Tables** to include / exclude history tables in the table. Right-click on a table cell to **Export to Excel** or **Copy Cell Value**.

This low-level option can be useful to identify the source of the problem when the displays are not showing the expected data. Use this data for debugging and troubleshooting with Technical Support.

Cache Table 07-May-2019 14:11 ✓ DATA

Data Server: central-alert History Tables: ☐

Data Server URL: https://rtvdemos.sl.com/emdemo_central_rtquery

Cache	Table	Rows	Columns	Memory
JmxStatsTotals	current	1	4	441
RtvAlertGroupMap	current	493	3	67424
RtvAlertMapByCI	current	62	5	13614
RtvAlertSourceStats	current	8	2	940
RtvAlertStatsByArea	current	8	9	2930
RtvAlertStatsByAreaAndAlertGroup	current	8	10	3454
RtvAlertStatsByCI	current	59	5	9228
RtvAlertStatsByCIAndAlertGroup	current	59	8	12506

Cache: **RtvAlertStatsByCIAndAlertGroup** Table: **current**

time_stamp	CITYPE	CINAME	ALERTGROUP	MaxSeverity	AlertCount
2019-May-07 14:11:33	JVM	localhost:SOLMON_TOM	None	2	1
2019-May-07 14:11:33	JVM	localhost:EMSMON_TOI	None	2	1
2019-May-07 14:11:33	JVM	localhost:EMSMON_DAT	None	2	1
2019-May-07 14:11:33	JVM	localhost:SOLMON_DISF	None	2	1
2019-May-07 14:11:33	JVM	localhost:SOLMON_DATI	None	2	1
2019-May-07 14:11:33	JVM	localhost:EMSMON_DISI	None	2	1
2019-May-07 14:11:33	JVM	localhost:SOLMON_TOM	None	2	1
2019-May-07 14:11:33	JVM	localhost:EMSMON_DAT	None	2	1
2019-May-07 14:11:33	JVM	Instance-1-90;CRMBroke	None	1	1
2019-May-07 14:11:33	JVM	Instance-1-90;CRMZooki	None	1	1
2019-May-07 14:11:33	JVM	Instance-1-171;CRMCon	None	1	1
2019-May-07 14:11:33	JVM	Instance-1-171;CRMCon	None	1	1
2019-May-07 14:11:33	JVM	Instance-1-171;CRMBrok	None	1	1
2019-May-07 14:11:33	JVM	localhost:TMolecule5_2	None	1	1
2019-May-07 14:11:33	JVM	localhost:PMolecule12_1	None	1	1

Page 1 of 2 1 - 40 of 59 items

Modify RTView Manager Settings

RTView Manager has predefined connections that can be modified if necessary. You can modify RTView Manager settings using the RTView Configuration Application here (<http://localhost:3070/rtview-manager-rtvadmin> or <http://localhost:8068/rtview-manager-rtvadmin>).

You can use the RTView Configuration Application to modify the majority of settings for RTView Manager.

This section contains:

- [“Open the RTView Configuration Application for RTView Manager”](#)
- [“Modify Connections for Data Collection”](#)
- [“Modify Default Polling Rates for RTView Manager Caches”](#)
- [“Modify Default Settings for Storing Historical Data”](#)
- [“Change Port Assignments”](#)
- [“Configure Alert & Historical Database Connections”](#)


Open the RTView Configuration Application for RTView Manager

To access the RTView Configuration Application for RTView Manager:

1. Start the Solace PubSub+ Monitor (if not currently running), then browse to one of the following URLs and login (username/password are rtvadmin/rtvadmin):
 - http://<ip_address>:3070/rtview-manager-rtvadmin if you are using Jetty.
 - <http://localhost:8068/rtview-manager-rtvadmin> if you are using Tomcat.

The RTView Manager main console opens.


The RTView Configuration Application **HOME** page opens.



Select the **RTView Server - RTView Manager Monitor** project.

The main configuration page for **RTView Manager** opens.

The navigation tree is in the left panel and the **General** and **Custom Properties** tabs are shown in the upper part of the main page. The name of the selected tab is highlighted and the other tabs are grayed out. You click on either of the grayed tabs to change the selected tab.




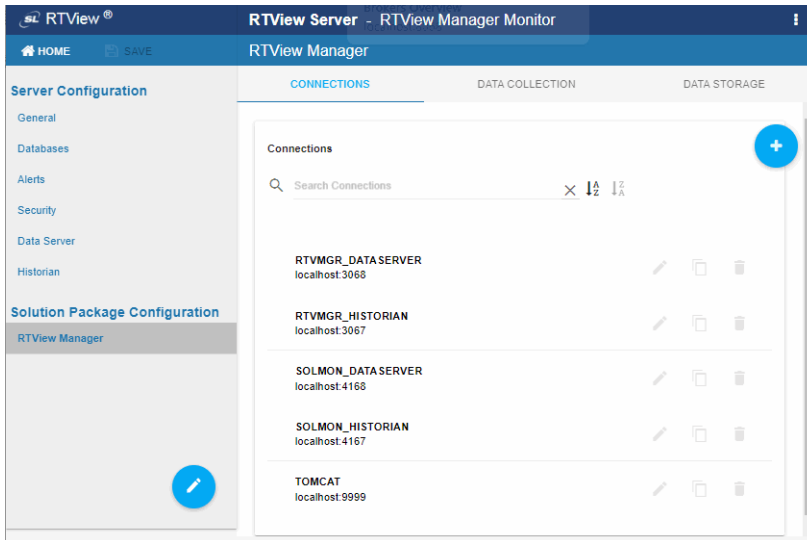
These instructions use the following format to describe navigation to each tab: **Navigation tree>Tab**. For example, the figure above illustrates the **General>GENERAL** Tab.

Modify Connections for Data Collection





RTView Manager has predefined connections to the Solace PubSub+ Monitor components. These connections do not need to be modified unless you change the ports or security settings for the Solace PubSub+ Monitor components

To modify connections:

1. “Open the RTView Configuration Application for RTView Manager”, select **RTView Manager>CONNECTIONS** tab and click .



The **Add Connection** dialog opens. Note that the predefined connections are listed in the main panel.

2. In the **Add Connection** dialog, enter the **Connection Name**, **Host**, **Port**, **Username** and **Password**.
3. Click  to close the dialog and  (in title bar) to save your settings.
The connections you create are listed in the **Connections** tab.
4. If your connection is secured, select **Security** (in the navigation tree) and fill in the **SSL Credentials** section with the appropriate **Truststore** and **Truststore Password** values for the connection.
5. Repeat these steps to add more brokers and when finished, click  to close the dialog and  (in title bar) to save your settings. RTView Manager

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

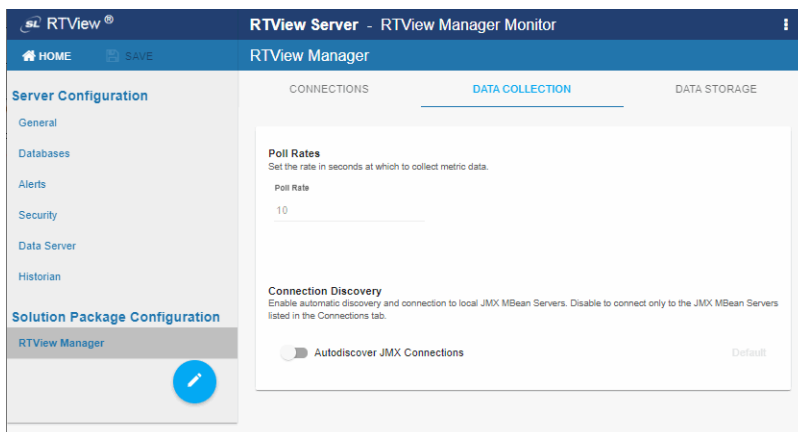
6. Click **RESTART SERVERS** to apply changes. The data server will be available again in 10-15 seconds.
7. Open a browser and go to the RTView Manager (Solace PubSub+ Monitor must be running):
 - **http://<ip_address>:3070/rtview-manager** if you are using Jetty.
 - **http://localhost:8068/rtview-manager** if you are using Tomcat.
 (username/password are rtvadmin/rtvadmin)

You should now see monitoring data for the modified connection. If you encounter issues, check the log files in the **projects/rtview-manager/logs** directory for errors.

Modify Default Polling Rates for RTView Manager Caches

To modify the default polling rate settings for RTView Manager caches, perform the following:

- “[Open the RTView Configuration Application for RTView Manager](#)” and go to **RTView Manager Monitor project>DATA COLLECTION** tab.



Poll Rates: Collection period in seconds. The default setting is **10** seconds.

Autodiscover JMX Connections: Toggle **ON** to enable RTView to automatically discover JMX MBean Servers. Toggle **OFF** to restrict connections to the JMX MBean Servers that are listed in the **CONNECTIONS** tab. Blue (toggled right) enables, gray (toggled left) disables. By default, this feature is disabled.

- Click **SAVE** your settings, then click **RESTART SERVERS** to apply changes. The data server will be available again in 10-15 seconds.

Modify Default Settings for Storing Historical Data

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

- “[Define the Storage of In Memory History](#)”: Specify the maximum number of history rows to store in memory.

- [“Define Compaction Rules”](#): Define rules for reducing the amount of data stored over time.
- [“Define Duration”](#): Specify when data becomes expired and/or deleted from the Monitor.
- [“Enable/Disable Storage of Historical Data”](#): Choose the metrics you want to store in the database and specify a prefix for history table names.
- [“Define Prefix for All History Table Names”](#): Specify a prefix to prepend to database table names.

Define the Storage of In Memory History



You can define the maximum number of history rows to store in memory in the **RTView Manager/Data Storage/History Rows** property. This property can improve Monitor responsiveness.

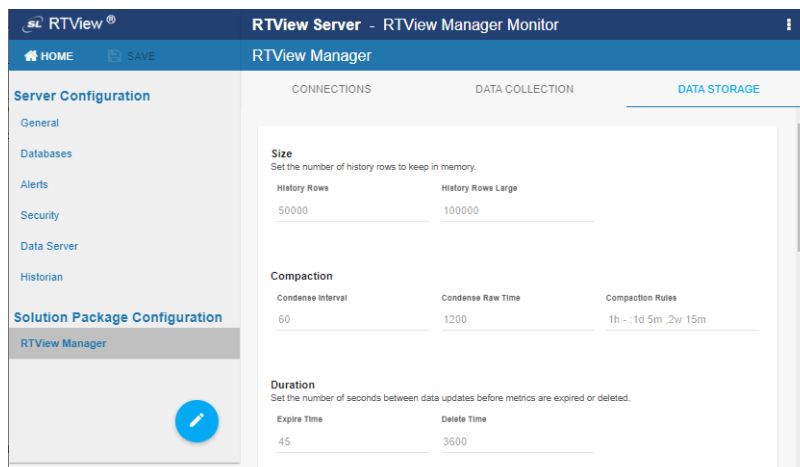
Note that changing this value is only recommended if you have a high degree of understanding about how historical data is being stored in memory, as well as how that data is compacted and stored in the database.

The **History Rows** property defines the maximum number of rows to store for the `JvmGcInfo`, `JvmMemoryPool`, `RtvDataServerManager`, `RtvDisplayServerManager` and `RtvDataServerClientTotals` caches. The default setting for **History Rows** is **50,000**.

The **History Rows Large** property defines the maximum number of rows to store for the `JvmOperatingSystem`, `JvmThreading`, `JvmMemory`, `RtvDataServerClientStats` and `TomcatWebModuleStats` caches. The default setting for **History Rows Large** is **100,000**.

To modify the default settings:

- [“Open the RTView Configuration Application for RTView Manager”](#) and go to **RTView Manager>DATA STORAGE** tab.
- Under **Size**, enter the desired number of rows in the **History Rows** and **History Rows Large** fields.
-  your settings, then click  to apply changes. The data server will be available again in 10-15 seconds.





Define Compaction Rules

Data compaction, essentially, is reducing redundancy in the data to be stored in the database by using a rule so that you store sampled data instead of raw data, which prevents storing of redundant data which potentially can overload the database. The compaction rule is defined through the following fields:

- **Condense Interval:** The time interval at which the cache history is condensed. The default is **60** seconds, which means that every **60** seconds all rows of the same index are condensed. As a result of this first condensing operation there will be only one row per index every minute. The following caches are impacted by this setting: `JvmGcInfo`, `JvmMemoryPool`, `JvmOperatingSystem`, `JvmThreading`, `JvmMemory`, `RtvDataServerManager` and `RtvDataServerClientTotals`.
- **Condense Raw Time:** The time span of raw data kept in memory. The default is **1200** seconds. The following caches are impacted by this setting: `JvmGcInfo`, `JvmMemoryPool`, `JvmOperatingSystem`, `JvmThreading`, `JvmMemory`, `RtvDataServerManager`, `RtvDataServerClientTotals`, `TomcatWebModuleStats`, `TomcatGlobalRequestStats` and `TomcatWebModuleTotals`.
- **Compaction Rules:** This field defines the rules used to condense your historical data in the database. By default, the columns kept in history are aggregated by averaging rows with the following rule **1h -;1d 5m;2w 15m**, which means the data from the last hour is not aggregated (1h - rule), the data from the last day is aggregated every 5 minutes (1d 5m rule), and the data from the last 2 weeks old is aggregated every 15 minutes (2w 15m rule). The following caches are impacted by this setting: `JvmOperatingSystem`, `JvmThreading`, `JvmMemory`, `RtvDataServerManager`, `RtvDataServerClientTotals`, `TomcatWebModuleStats`, `TomcatGlobalRequestStats` and `TomcatWebModuleTotals`.

To modify these settings do the following:

- [“Open the RTView Configuration Application for RTView Manager”](#) and go to **RTView Manager>DATA STORAGE** tab.
- Under **Compaction**, enter values in the **Condense Interval**, **Condense Raw Time** and **Compaction Rules** fields.
-  **SAVE** your settings, then click  **RESTART SERVERS** to apply changes.



Define Duration

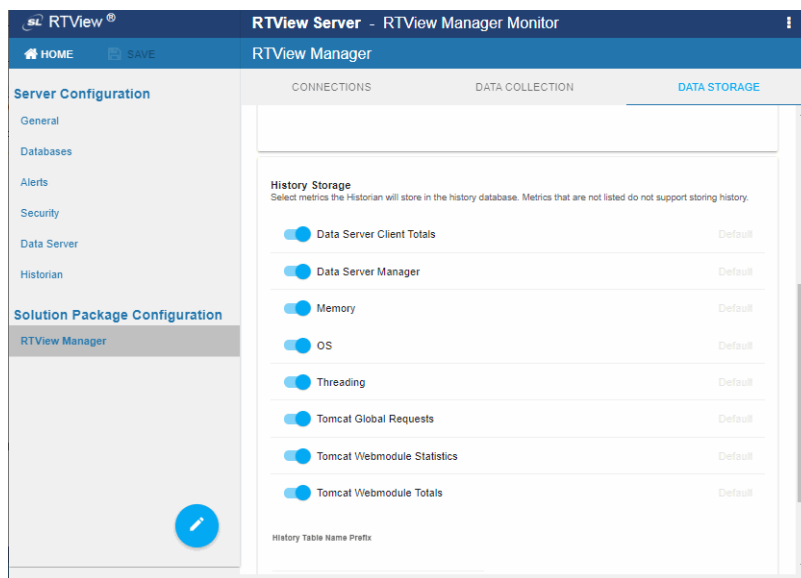
The data for each metric is stored in a specific cache and, when the data is not updated in a certain period of time, that data either marked as expired or, if it has been expired over an extended period of time, it is deleted from the cache altogether.

- **Expire Time:** This field sets the period of time when the **Expire** metric from the cache is set to true indicating the entry row is expired. The default expiration time is **45** seconds. The following caches are impacted by this field: JvmConnections, JvmGcInfo, JvmMemoryPool, JvmClassLoading, JvmCompilation, JvmOperatingSystem, JvmThreading, JvmMemory, JvmMemoryManager, JvmSystemProperties, RtvDataServerManager, RtvDisplayServerManager, RtvHistorianManager, RtvDataServerClientStats, RtvDataServerClientTotals, RtvServerVersion, TomcatWebModuleStats, TomcatConnectorInfo, TomcatGlobalRequestStats, TomcatHostInfo, and TomcatWebModuleTotals.
- **Delete Time:** This field sets the period of time that a given entry row should be expired before it gets deleted from the cache. It defaults to **3600** seconds and applies to the following caches: JvmConnections, JvmGcInfo, JvmMemoryPool, JvmClassLoading, JvmCompilation, JvmOperatingSystem, JvmRuntime, JvmThreading, JvmMemory, JvmMemoryManager, JvmSystemProperties, RtvDataServerManager, RtvDisplayServerManager, TomcatWebModuleStats, TomcatGlobalRequestStats, TomcatWebModuleTotals, RtvHistorianManager, RtvDataServerClientStats, RtvDataServerClientTotals, RtvServerVersion, TomcatWebModuleStats, TomcatConnectorInfo, TomcatGlobalRequestStats, TomcatHostInfo and TomcatWebModuleTotals.

Enable/Disable Storage of Historical Data

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:

- “Open the [RTView Configuration Application for RTView Manager](#)” and go to **RTView Manager>DATA STORAGE** tab.
- 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 Brokers), SolBridgeStats (Bridge Stats), SolClientStats (Client Stats), SolCspfNeighbors (CSPF Neighbors), SolEndpointStats (Endpoint Stats), SolEndpoints (Endpoints), SolApplianceInterfaces (Interface), SolApplianceMessageSpool (Message Spools), SolEventModuleEvents (Syslog Events) and SolVpns (VPNs).
-  **SAVE** your settings, then click  **RESTART SERVERS** to apply changes.





Define Prefix for All History Table Names

The **History Table Name Prefix** field allows you to define a prefix that is 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 need to create the corresponding tables in your database as follows:

- Locate the .sql template for your database under `/rtvapm/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 Configuration Application for RTView Manager”, go to **RTView Manager>DATA STORAGE** tab and scroll down to the bottom of the page.
- In the **History Table Name Prefix** field, enter the desired prefix name.
-  your settings, then click  to apply changes.

Change Port Assignments

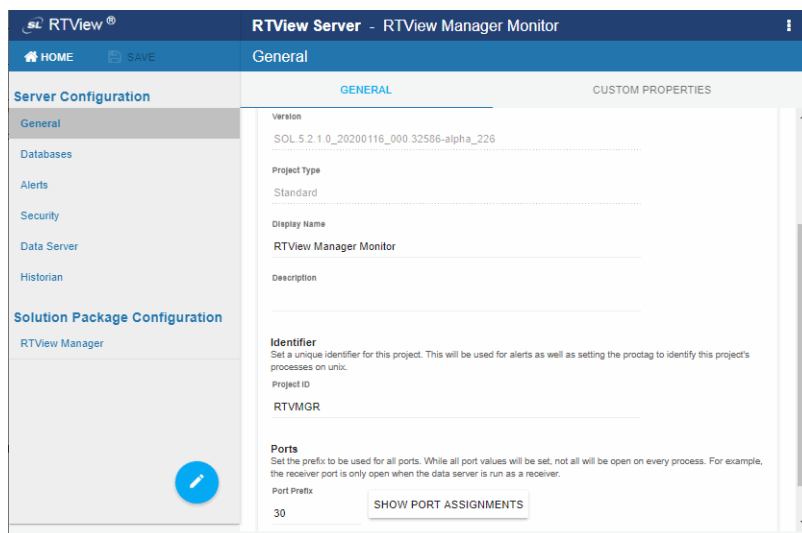
This configuration is optional.

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 RTView Manager using the RTView Configuration Application.



Java Process	Description	Default Port(s)
RTView Manager Data Server	Gathers performance metrics.	Default Port= 3078 Default JMX Port = 3068
RTView Manager Historian	Retrieves data from the RTView Data Server and archives metric history to a database.	Default JMX Port= 3067

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

1. “Open the RTView Configuration Application for RTView Manager” and go to **General>GENERAL** tab.



2. Under **Ports** (scroll down to the bottom of the page), 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** your settings, then click  **RESTART SERVERS** to apply changes.
4. Edit the **update_wars** (.bat or .sh) file and change the port prefix for all ports to the prefix you just specified.
5. Rebuild the war files and install them to the application server by executing the following script, located in the **/bin** directory:

Windows:

make_all.bat

UNIX:

./make_all.sh

Configure Alert & Historical Database Connections

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*.

This section describes how to setup an alternate production database, and how to configure the Alert Settings Database connection and the Historian Database connection. You connect and configure the databases using the RTView Configuration Application. You also copy portions of the **database.properties** template file (located in the **common\dbconfig** directory) into the RTView Configuration Application.

Monitor Databases

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 cannot work correctly.

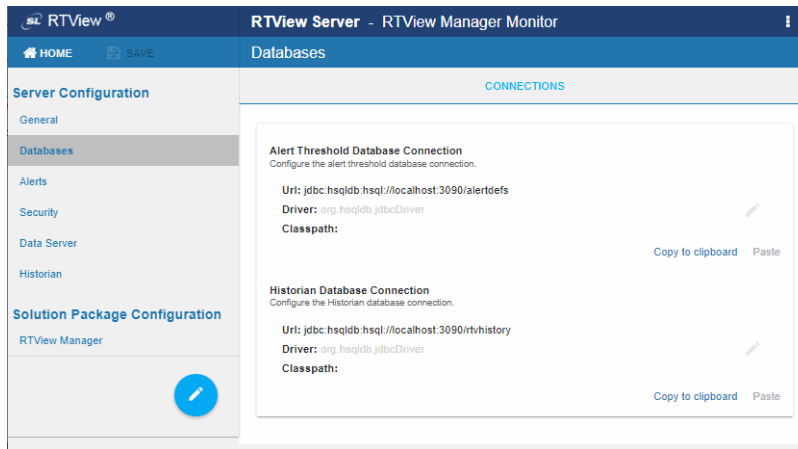
Historian

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 ALERTDEFS and RTVHISTORY Databases:

1. Install a database engine of your choice. Supported database engines are Oracle, Microsoft SQL Server, MySQL, and DB2.
NOTE: The default page size of DB2 is 4k. It is required that you create a DB2 database with a page size of 8k. Otherwise, table indexes will not work.
2. Open the **database.properties** template file, which is located in the **common\dbconfig** directory, find the line that corresponds to your supported database in the "Define the ALERTDEFS DB" section and make a note of this information. Keep the **database.properties** template file open.

3. “Open the RTView Configuration Application for RTView Manager” and go to **Databases>CONNECTIONS** tab.



4. Click **Alert Threshold Database Connection** to open the **Edit Connection** dialog.
5. Enter the information (you previously noted from the **database.properties** file) 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.

Click **SAVE** to close the dialog and **SAVE** (in title bar) to save your settings.
6. Return to the **database.properties** template file, which is located in the **common\dbconfig** directory, find the line that corresponds to your supported database in the "Define the RTVHISTORY DB" section and make a note of this information.
7. In the RTView Configuration Application, click the **Historian Database Connection** to open the **Edit Connection** dialog.
8. Enter the information (you previously retrieved from the **database.properties** file) 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.

9. Click  to store the newly added connection and close the dialog and  (in title bar) to save your settings.

10. Click  to apply changes.

11. 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, where:

`<db> = {db2, mysql, oracle, sqlserver}`

- **Alert Settings**

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

- **Historian**

`SolacePubSubMonitor/rtvapm/common/dbconfig/
create_common_history_tables_<db>.sql`

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

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. To configure alert notifications, proceed to [Configure Alert Notification](#).

Troubleshoot

This section includes:

- [“Log Files for RTView Manager”](#)
- [“JAVA_HOME”](#)
- [“Permissions”](#)
- [“Network/DNS”](#)
- [“Data Not Received from Data Server”](#)

Log Files for RTView Manager

When any 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.

RTView Manager Log Files

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

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

JAVA_HOME

If you encounter issues starting Solution Package for Solace 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 confirm with your network administrator that you're not being blocked from accessing the remote system.

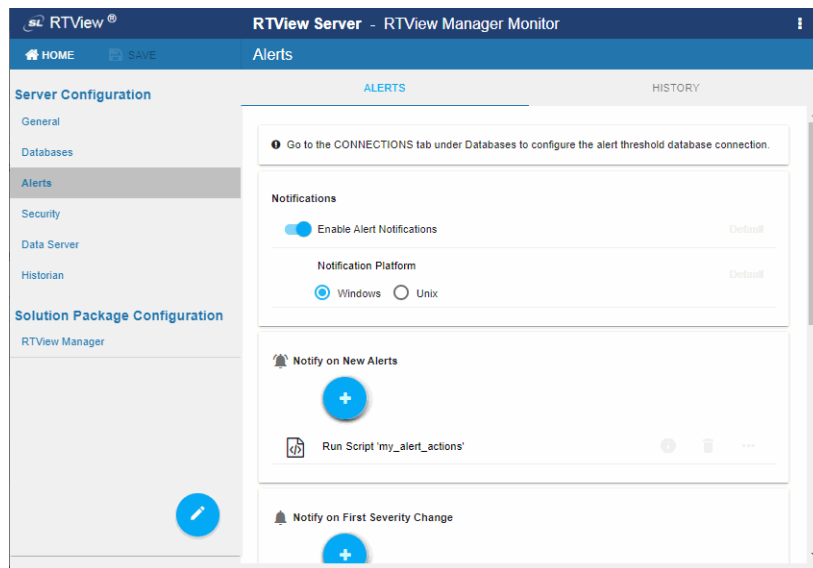
Data Not Received from Data Server

In the Solution Package 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. Do the following:

1. [“Open the RTView Configuration Application for RTView Manager”](#) and go to the **RTView Manager>CONNECTIONS** tab.
2. Verify the settings for each connection and make corrections if necessary.
Click **SAVE** in the title bar when finished, then click **RESTART SERVERS** to apply changes. It takes about 10-15 seconds for the data server to be available again.
3. In Solace PubSub+ Monitor, go to the **Admin>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).

Configure Alert Notification

To configure alert notification for RTView Manager, [“Open the RTView Configuration Application for RTView Manager”](#), select the **RTView Manager** project, select **Alerts** (in the navigation tree) and follow [“Configure Alert Notification”](#) instructions.



For details about alerts for RTView Manager, see [“Alerts for RTView Manager”](#).

Configure High Availability

To configure HA for RTView Manager, refer to [“High Availability”](#) instructions.

APPENDIX A RTView Configuration Application for Solace PubSub+ Monitor

This section describes the RTView Configuration Application pages and settings for configuring the Solace PubSub+ Monitor *projects*.

The “Overview” provides general information about:

- “Open the RTView Configuration Application” via browser or from the Monitor.
- “The HOME Page”: where you select the project you want to configure.

The screenshot shows the RTView Configuration Application interface. At the top, there is a blue header bar with 'HOME' and 'SAVE' icons on the left, and 'General' in the center. Below the header, the main content area is divided into two sections: 'Server Configuration' on the left and 'GENERAL' on the right. The 'Server Configuration' section has a sidebar with links: 'General', 'Databases', 'Alerts', 'Security', 'Data Server', 'Historian', 'Solution Package Configuration', 'Data Cache', and 'Solace'. The 'GENERAL' section contains the following fields:

- About**
 - URL: localhost:4178
 - Location: /home/m/testbed/TBSOLMON/SolacePubSubMonitor/projects/rtview-server
 - Version: SOL 5.3.0.0_20200721_000.33415-alpha_145
 - Project Type: Standard
 - Display Name: Solace Monitor
 - Description: (empty field)
- Identifier**
 - Set a unique identifier for this project. This will be used for alerts as well as setting the proctag to identify this project's processes on unix.
 - Project ID: SOLMON
- Ports**
 - Set the prefix to be used for all ports. While all port values will be set, not all will be open on every process. For example, the receiver port is only open when the data server is run as a receiver.
 - Port Prefix: 41
 - SHOW PORT ASSIGNMENTS button

At the bottom of the sidebar, there is a blue circular button with a white pencil icon.

This section then describes each page in the order that they appear in the navigation tree (the left panel in the figure above).


- **"General Page"**: Provides details about your Monitor installation (under the **"GENERAL Tab"**) including current **Port Assignments**, and allows you to enter and modify custom properties (under the **"CUSTOM PROPERTIES Tab"**).
- **"Databases Page"**: Describes the page in which you setup connections for the Alert Threshold Database and the Historian Database.
- **"Alerts Page"**: Describes the pages in which you configure alert notifications, alert persistence (**"ALERTS Tab"**) and alert storage (**"HISTORY Tab"**).
- **"Security Page"**: Describes the page in which you secure JMX ports with SSL and username and password.
- **"Data Server Page"**: Describes the pages in which you set the initial and maximum amount of memory for Data Server processing, log file name and location, set the Data Server to host servlets (**"DATA SERVER Tab"**) and add data server connections (**"COLLECTOR Tab"**).
- **"Historian Page"**: Describes the page in which you allocate memory and set log files for the Historian.
- **"Solution Package Configuration"**: Describes the page in which you setup a Solution Package.
- **"Custom Display Designer"**: Describes how to quickly create and deploy custom displays.

Overview

After you log in the **HOME** page opens. This is where you select the project you want to configure.

Open the RTView Configuration Application

1. In the Solace PubSub+ 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**

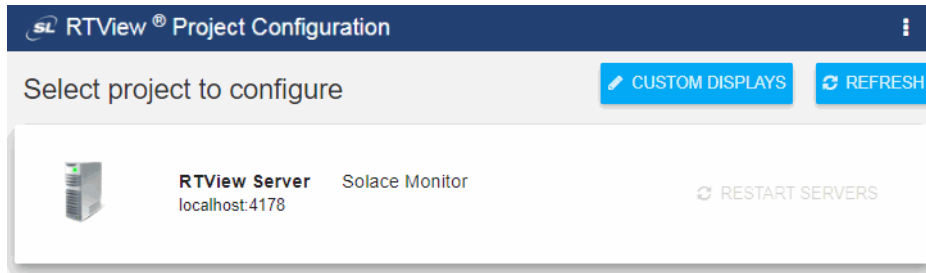
Password: **rtvadmin**

The RTView Configuration Application **HOME** page opens.

3. Select the project you want to configure (scroll down if necessary).

The HOME Page

The home page provides access to your Solace Monitor project (as well as your Solution Package Projects to which the Monitor is connected). The Solace Monitor project allows you to configure alerts, databases, data servers, the historian, the display server as well as data collection and rules for data storage.




The **CUSTOM DISPLAYS** button is a shortcut to the [“Custom Display Designer”](#).

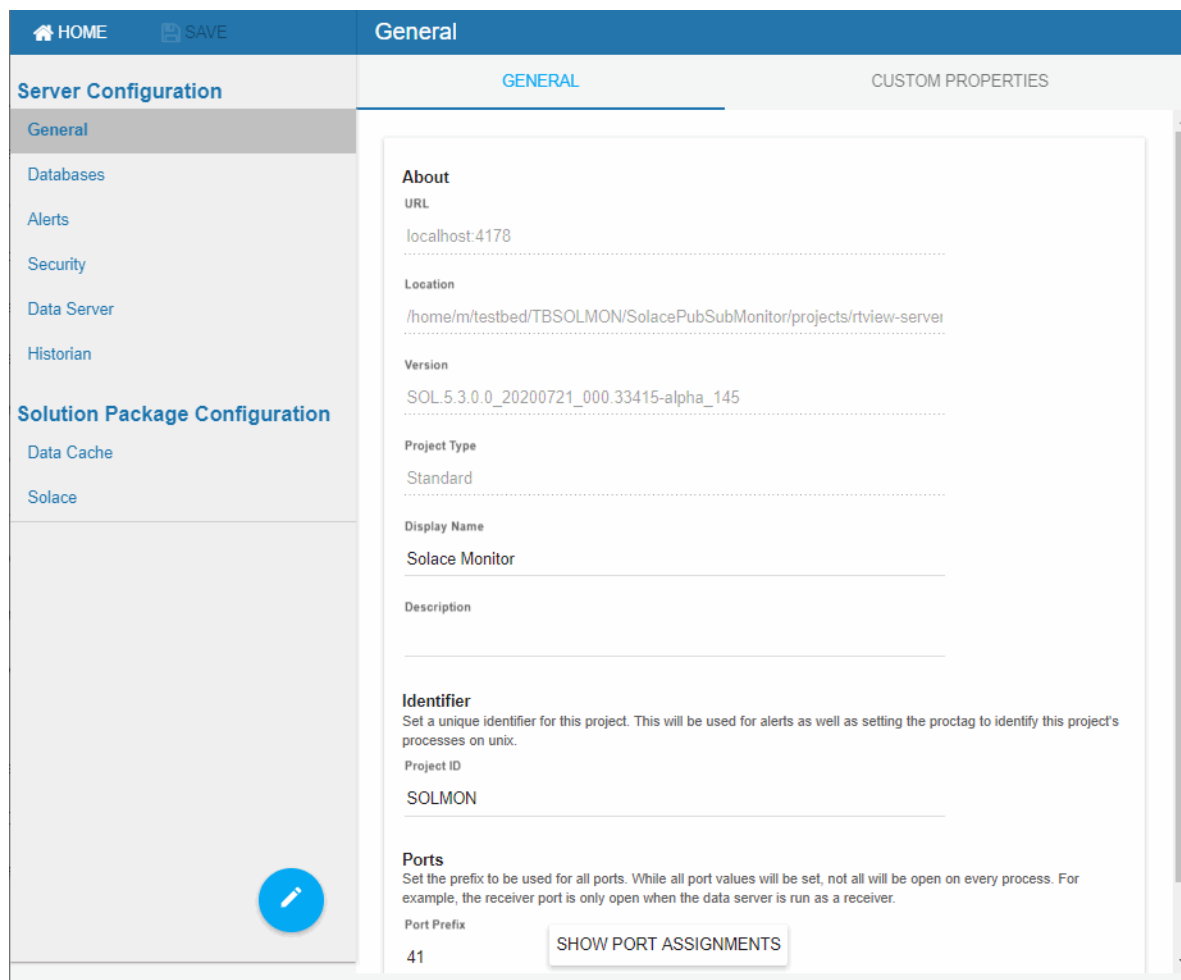
The **REFRESH** button refreshes the project list.

Select a project in the list to configure.


The [“General Page”](#) is the main configuration page and opens by default.


Main Configuration Application Page

The “General Page” opens by default. To return to the HOME page click  (in the title bar).





Saving and Applying Settings



 is located in dialogs in which you enter new settings. This button closes the dialog and saves recent settings in memory.

 is located in the title bar. This button saves recent settings locally.


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):

 is visible when you have an unapplied changes.

Click  to see the Solution Packages that are installed, select those you want to add.

Click  to close the dialog and .

To apply changes to your system click [RESTART DATASERVER](#).



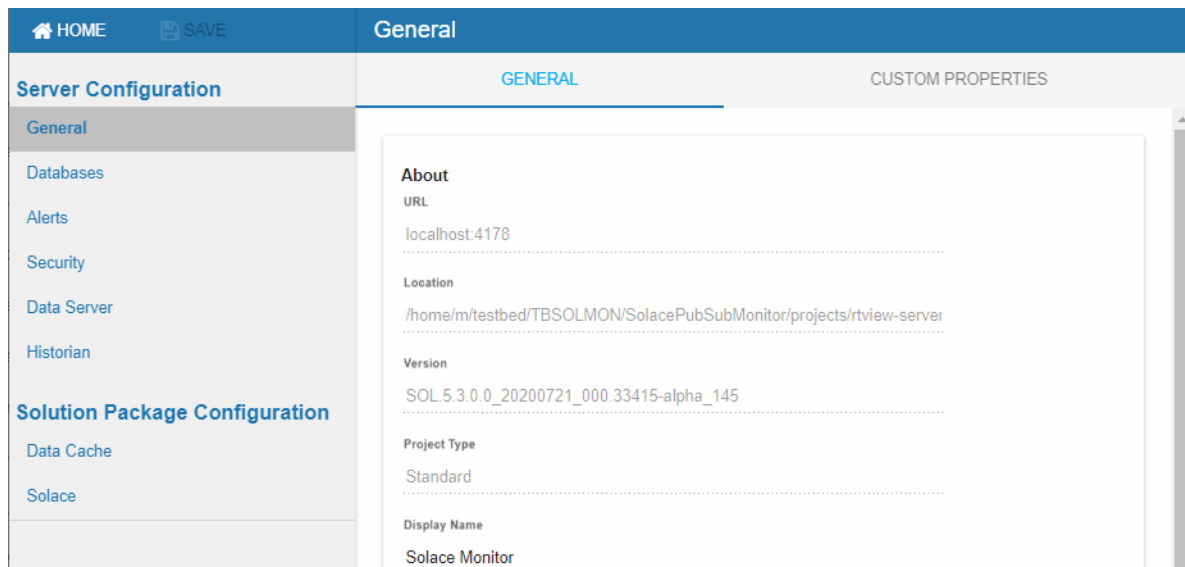
The Solution Packages immediately appear in the navigation tree under **Solution Package Configuration**.

You must restart the server before you can configure the Solution Package.

General Page

This page has two tabs:

- **"GENERAL Tab"**: Get details about your Monitor installation from this page.
- **"CUSTOM PROPERTIES Tab"**: Use this page to enter custom properties.



GENERAL Tab

The **GENERAL** tab (shown above) provides the following details about your Monitor installation:

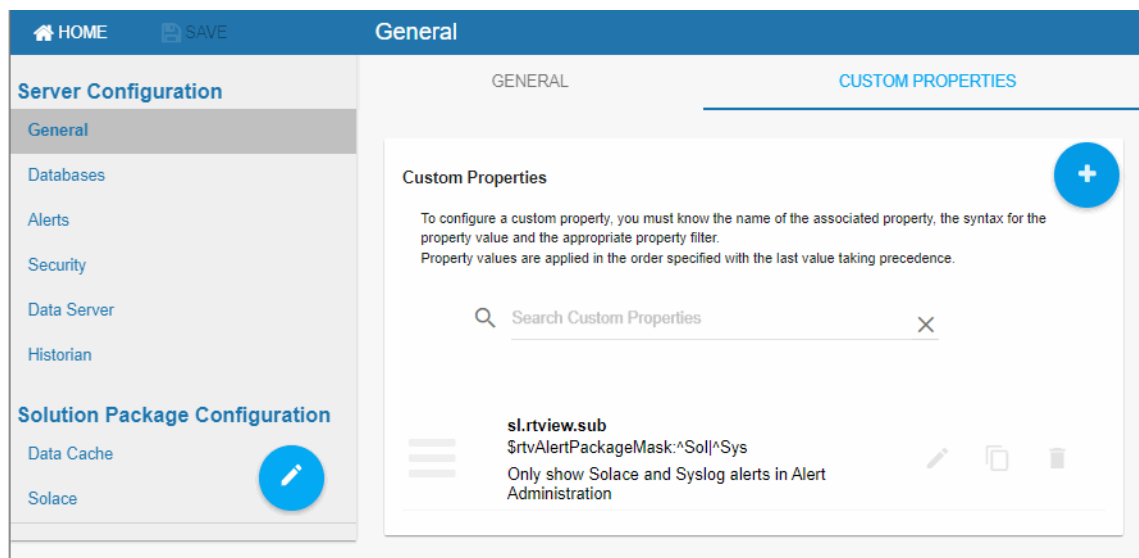
Field Name	Description
Display Name	The name of the Monitor display.

Description	Administrator's notes.
Identifier	The unique identifier for this project which is used for alerts and UNIX processes.
Port Prefix	The port prefix to use for all ports.
Show Port Assignments	Click to see list of port assignments.

CUSTOM PROPERTIES Tab


Use the **CUSTOM PROPERTIES** page to create and search for custom properties for the Server.

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.



Click  to add a custom property in the **Add Property** dialog.

To edit, clone or delete a property, hover your mouse over the property to bring the option icons into focus. For example:



The **Add Property** dialog has the following fields:

Field Name	Description
Name	The property name.
Value	The property value.

Filter	The property filter (optional).
Comment	A comment describing this property (optional).

Databases Page

Use the **Databases** page to setup the Alert Threshold Database Connection and the Historian Database Connection.

Click **Copy to clipboard** on any RTView Servers database field to copy it to the clipboard. Then click **Paste** on any of the other connection listed.

The settings you make here are applied to all Monitor servers. For example, when you change the historian database connection in this tab, the database connection for the Alert Historian also changes, as well as for the Solution Package for Solace servers.

The **Databases** page has the following fields:

Field Name	Description
Alert Threshold Database Connection	<p>This is the connection to use for the Alert Threshold database. This database contains all alert settings (warning and alarm thresholds, etc).</p> <p>URL: Full URL to use when connecting to this database using the specified JDBC driver.</p> <p>Driver: Fully qualified name of the driver class to use when connection to this database via JDBC.</p> <p>Classpath: The classpath to the jar containing the driver class.</p> <p>Username: (optional) User name to enter into this database when making a connection.</p>

Password: (optional) Password to enter into this database when making a connection.

Click **Create new password** to change your password.

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.

Historian Database Connection

This is the connection to use for the Historian database.

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.

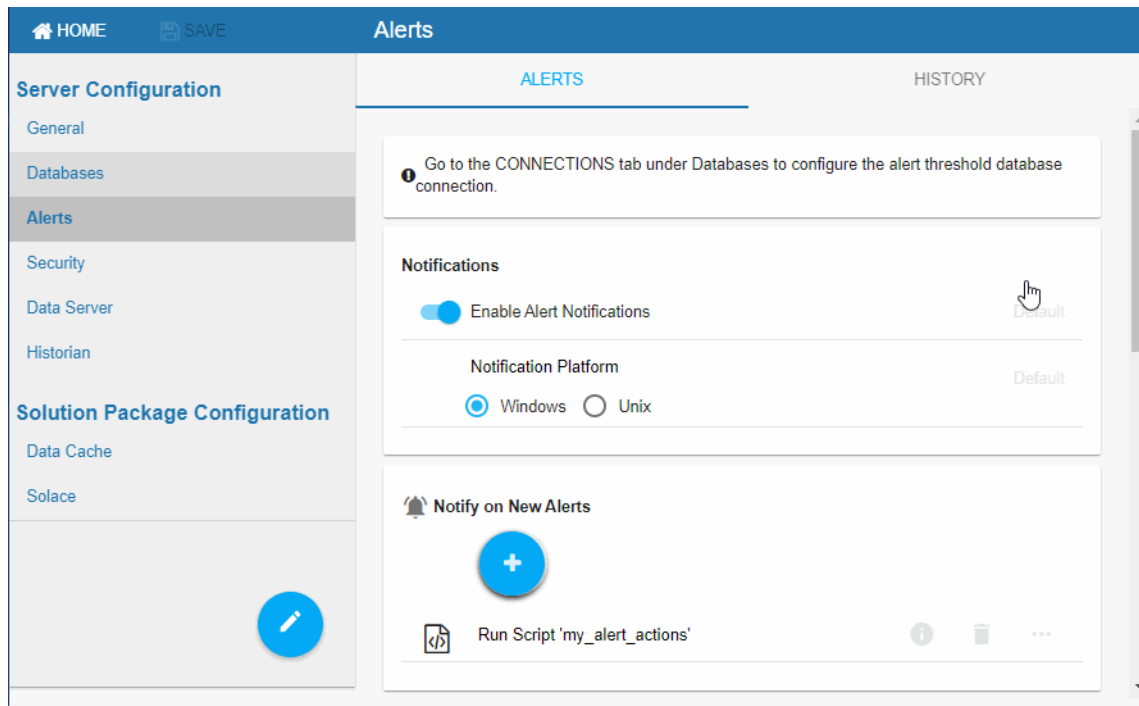
Alerts Page

Use the **Alerts** page to configure alert notifications, alert persistence for high availability and alert storage. The Alerts page has two tabs:

- **"ALERTS Tab":** Configure alert notifications and alert persistence for high availability.
- **"HISTORY Tab":** Configure alert storage.

ALERTS Tab

Use the **ALERTS** tab to configure alert notifications and persistence for high availability.



The **ALERTS** tab has the following fields:

Field Name	Description
Enable Alert Notifications	Toggle on to enable alert notifications. By default, alert notifications will execute a script in the servers/central directory.
Notification Platform	Select the platform of the system that the alert server is running on.
Notify on New Alerts	Toggle on 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	Toggle on 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/central and modify it to execute the action you want to perform.
Notify on Cleared Alerts	Toggle on 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	Toggle on 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, in seconds, on which you want to renotify on unacknowledged alerts. The default is **300** seconds.

Persist Alerts

Toggle on to persist the current alert table to the Alert Threshold Database.

HISTORY Tab

Use the **HISTORY** tab to configure alert storage in the history database.

The screenshot shows the RTView Configuration Application interface. On the left is a sidebar with navigation links: HOME, SAVE, Server Configuration (General, Databases, Alerts, Security, Data Server, Historian), Solution Package Configuration (Data Cache, Solace), and a blue circular icon with a pencil. The main area is titled 'Alerts' and has two tabs: 'ALERTS' and 'HISTORY'. The 'HISTORY' tab is active. It contains a 'History' section with the text 'Select to have the Historian save alerts to the history database.' Below this is a toggle switch for 'Store Alert History', which is currently turned off. To the right of the toggle is the word 'Default'. Below the toggle is a text input field for 'History Table Name Prefix'. Underneath the input field is a detailed instruction: 'Enter a value to prepend to the history table names for all metrics. Note that this requires a change to your history database schema. If you are using Oracle for your Historian Database, you must limit the History Table Name Prefix to 2 characters because Oracle does not allow table names greater than 30 characters.'

The **HISTORY** tab has the following fields:

Field Name**Description****Store Alert History**

Toggle on to have the 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 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.

NOTE: If you are using Oracle for your Historian Database, you must limit the History Table Name Prefix to 2 characters because Oracle does not allow table names greater than 30 characters.

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.

Security Page

Use the **Security** page to secure RTView JMX ports.

All RTView processes (Data Server, Historian, Display Server) open JMX ports for monitoring which, by default, are not secured. The Security tab allows you secure these ports as well as specify credentials needed to connect to SSL secured servers from RTView's Solution Packages.

SSL Credentials

This region allows you to specify the path to the Truststore and Keystore files (and their associated passwords) that contain the SSL credentials needed to secure the RTView JMX Ports and/or access any SSL secured servers associated with RTView's Solution Packages. This is required if the Secure with SSL option is enabled (see below for details).

Securing RTView JMX Ports

This region provides a couple of options for securing the JMX ports that are opened by the RTView processes: Secure with SSL and/or Secure with Username and Password.

Secure with SSL

When toggled on, this option secures the JMX ports for the RTView processes with SSL. When the JMX ports are SSL secured, an SSL handshake is performed between the client and the server when the client attempts to connect. For this handshake, the client must provide a certificate the server trusts, and the server must provide a certificate the client trusts. A Keystore file contains the application's certificate and private key and a Truststore file contains the application's trusted certificates. These files are created by running the Java keytool on the command line. When this option is enabled, you must specify the path to the server's Truststore and Keystore files (and their associated passwords) in the SSL Credentials region (see above).

The **start_server**, **stop_server**, and **status_server** scripts are all connected to the JMX Ports of the RTView processes to execute operations and get status. If the JMX ports have been secured with SSL, these scripts need the path and passwords for the truststore and keystore files containing the client credentials in order to connect. You can either pass these in on the command line each time you run (**-sslkeystore:clientkeystore.jks -sslkeystorepass:clientkeystorepass -ssltruststore:clienttruststore.jks -ssltruststorepass:clienttruststorepass**) or you can fill in the fields under SSL Credentials for RTView Scripts.

The RTView Manager application also connects to the JMX Ports of the RTView processes in order to monitor them. If you are using the RTView Manager and the JMX ports have been secured with SSL, you must fill in the SSL Credentials on the Security tab of the RTView Manager Configuration Application to specify the path the truststore and keystore files containing the client credentials.

Secure with Username and Password

This region allows you to secure the JMX ports for RTView processes, which would then require using a specific username/password to gain access. If this option is enabled, you must specify the path to the password file containing all valid user names and passwords.

The **start_server**, **stop_server**, and **status_server** scripts are all connected to the JMX Ports of the RTView processes to execute operations and get status. If the JMX ports have been secured with a username and password, the scripts need a valid user name and password in order to connect. You can either pass these into the command line each time you run (**-jmxuser:userName -jmxpass:myPassword**) or you can fill in the Username and Password Credentials and enable the Use for Scripts toggle.

The RTView Manager application also connects to the JMX Ports of the RTView processes in order to monitor them. If you are using the RTView Manager in RTViewCentral and the JMX ports have been secured with username and password, you must fill in the Username and Password Credentials that the RTView Manager should use to connect. If you are using the RTView Manager in a deliverable other than RTViewCentral, you will need to fill in the user name and password in the connection to this RTViewDataServer in the RTView Manager Configuration Application.

The **Security** page has the following fields:

Option	Description
SSL Credentials	<p>Truststore: Enter the directory path to the truststore file and click the associated SET PASSWORD button to define the password required to access the file.</p> <p>Keystore: Enter the directory path to the keystore file and click the associated SET PASSWORD button to define the password required to access the file.</p>

Secure with SSL

Enabling this toggle locks the JMX ports for the RTView processes (DataServer and Historian, for example). You must specify the path to the Truststore and Keystore files (and their associated passwords) in the SSL Credentials region (see above) when selecting this toggle.

This option also locks the JMX ports used when the start_server, stop_server, and status_server scripts are run. Though not required, you can enter the path to the truststore and keystore files in the associated Client Truststore and Client Keystore fields (as well as their passwords) to avoid having to enter the keystore and truststore names and passwords on the command line when using the start_server, stop_server, and status_server scripts.

SSL Credentials for RTView Scripts: The start_server, stop_server, and status_server scripts connect to the RTView processes using JMX. You can either save the Client Truststore and Client Keystore properties here for use by the scripts or you can pass them in on the command line each time you execute those scripts.

For example, start_server.sh -sslkeystore:clientkeystore.jks -sslkeystorepass:clientkeystorepass -ssltruststore:clienttruststore.jks -ssltruststorepass:clienttruststorepass.

Secure RTView JMX Ports with Username and Password**Secure with User Name and Password**

Enabling this option secures the JMX ports for RTView processes, which would then require using a specific username/password to gain access. If this option is enabled, you must specify the path to the password file containing all valid user names and passwords.

Password File – enter the path to the password file containing the login and password credentials required to access the RTView processes.

Username and Password Credentials – If you are using RTView Manager in RTViewCentral and the JMX ports have been secured with username and password, you must fill in the Username and Password Credentials that the RTView Manager should use to connect.

Use for Scripts – Toggle on this option to allow the start_server, stop_server, and status_server scripts to use the username and password entered in the Username and Password Credentials region rather than having to specify the username and password on the command line each time you run the start_server, stop_server, and status_server scripts. For example, start_server.sh -jmxuser:userName -jmxpass:myPassword.

Data Server Page

Use the **Data Server** page to set the initial and maximum amount of memory for Data Server processing, log file name and location, set the Data Server to host RTView servlets, and add data server connections. The Data Server page has two tabs:

- **“DATA SERVER Tab”**: Configure initial and maximum amount of memory for Data Server processing, log file name and location and set the Data Server to host servlets.

- **"COLLECTOR Tab"**: Add, enable and setup sender and receiver Data Servers and set sender log file location.

The screenshot shows the 'Data Server' configuration page. The sidebar on the left contains navigation links: HOME, SAVE, Server Configuration (General, Databases, Alerts, Security, Data Server, Historian), and Solution Package Configuration (Data Cache, Solace). The main content area is titled 'Data Server' and has two tabs: 'DATA SERVER' (active) and 'COLLECTOR'. Under the 'DATA SERVER' tab, there are three sections: 'Memory' (Initial Memory: 256 MB, Max Memory: 4 GB), 'Logs' (Log File: logs/dataserver.log), and 'HTML Server' (HTML Server Enabled).

DATA SERVER Tab

The **DATA SERVER** tab (shown above) 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.
HTML Server Enabled	Toggle on to 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 automatically 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	Toggle on to enable 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.



COLLECTOR Tab

Use the **COLLECTOR** tab to setup the source and target data servers (the data servers that will send data and the data servers that will receive data). This tab is only visible if the project type is sender (for example, the project is run with the **-propfilter:sender** command line argument).

Click  to add a new Data Server target in the **"Add Sender Target Dialog"**, enter a unique ID and a URL, choose **All solution packages** to send data from all solution packages or choose **Select solution packages** to individually select from a list of Solution Packages.

To edit, copy or delete a Data Server, hover your mouse over the Data Server to bring the edit, copy and delete icons into focus. For example:





☒ Enable and  **SAVE** settings when finished. Then  **SAVE** (in title bar) to save changes locally.

The **DATA SERVER** tab has the following fields:

Field Name	Description
Sender Targets	This is a list of receiver Data Servers that the sender Data Server sends data to.
<input checked="" type="checkbox"/>	When checked, enables the Data Server to receive data from sender Data Server.
Log File	The log file name and location relative to the startup directory for the sender Data Server process.
Sender Identifier	The unique ID for the sender Data Server.

Add Sender Target Dialog

The **Add Sender Target** dialog has the following fields:

Field Name	Description
ID	The unique ID for the receiver Data Server.
URL	The URL for the receiver Data Server.
Targets	<p>All solution packages choose this option to send data from all solution packages to the receiver Data Server.</p> <p>Select solution packages choose this option to see a list of all available solution packages, and send data from a subset of those.</p>
<input checked="" type="checkbox"/>	When checked, enables the Data Server to receive data from sender Data Server.
	Saves the settings locally. To apply changes use  SAVE (in the title bar).

Historian Page

Use the **Historian** page to set the initial and maximum amount of memory for historian processing and set the log file name and location.

The **Historian** page has the following fields:

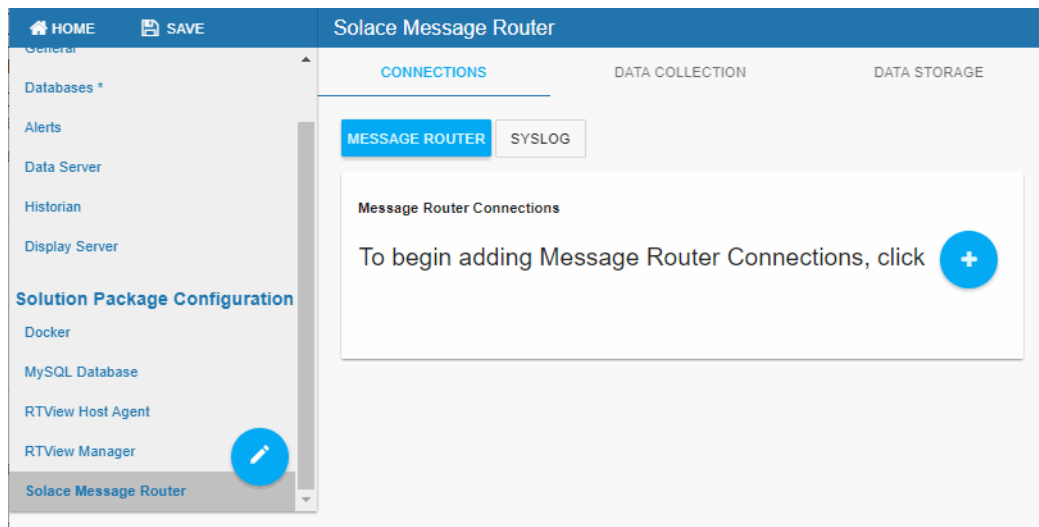
Field Name	Description
Initial Memory*	The initial amount of memory to allocate for this process.
Max Memory*	The maximum amount of memory to allocate for this process.
Log File	The log file name and location relative to the startup directory for this process.

***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.


Solution Package Configuration

This section describes the RTView Configuration Application page for configuring a Solution Package. There are three tabs:


- **"CONNECTIONS Tab"**: Set connections for Brokers.
- **"DATA COLLECTION"**: Set poll rates for Brokers.
- **"DATA STORAGE"**: Set how data is collected and aggregated, and reduce the amount stored.

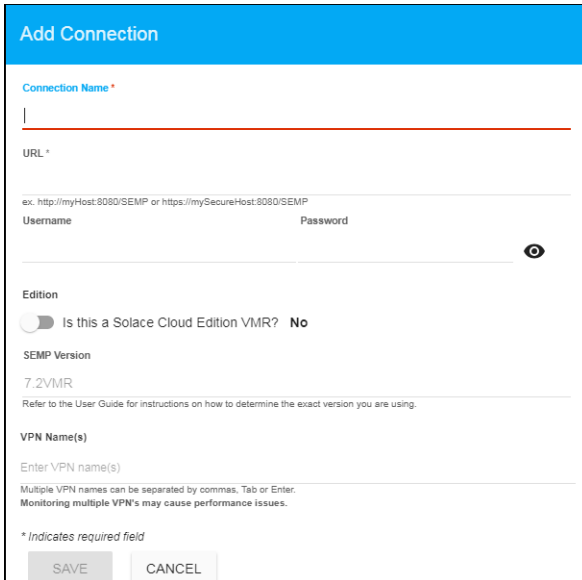


CONNECTIONS Tab

Use the CONNECTIONS tab to add connections. Select **"Broker"**, then click  to enter connection details in the **Add Connection** dialog.

Broker

Select MESSAGE ROUTER, then click  to add a broker connection in the **Add Connection** dialog.




The 'Add Connection' dialog box is shown. It has a blue header with the title 'Add Connection'. Below the header, there are several fields: 'Connection Name *' (required), 'URL *' (required), 'Username' (optional), 'Password' (optional, with a toggle icon), 'Edition' (toggle for 'Is this a Solace Cloud Edition VMR? No'), 'SEMP Version' (7.2VMR), and 'VPN Name(s)' (optional). A note at the bottom states: '* Indicates required field'. At the bottom right, there are 'SAVE' and 'CANCEL' buttons.

The MESSAGE ROUTER **Add Connection** dialog has the following fields:

URL: Enter the URL to be used for the connection, which contains an IP address or the host name that can be resolved by your network name resolution method and the SEMP port number configured for your broker.

Username: The username is used when creating the connection to the broker. This field is optional.

Password: This password is used when creating the connection to the broker. This field is optional. By default, the password entered is hidden. Click the  icon to view the password text.

Edition: Enable the toggle if the broker is a Solace Cloud Edition PubSub+.

SEMP Version: Specify the SEMP Version used by the broker.

VPN Name(s): Optionally enter the names of the VPNs that you want to monitor. Monitoring multiple VPNs could impact performance.

DATA COLLECTION

Use the **DATA COLLECTION** tab to modify the default values for data update frequency for various server-related caches, and also to specify **Connection Discovery** for automatic discovery and connection to local JMX MBean Servers.

The screenshot shows the RTView Server - Solace Monitor interface. The top navigation bar includes 'HOME', 'SAVE', and 'Solace'. The left sidebar lists 'Server Configuration' (General, Databases, Alerts, Security, Data Server, Historian) and 'Solution Package Configuration'. The main content area has three tabs: 'CONNECTIONS', 'DATA COLLECTION' (selected), and 'DATA STORAGE'. Under 'DATA COLLECTION', there are two sections: 'Poll Rates' and 'Event Module'. The 'Poll Rates' section has two input fields: 'Poll Rate' (set to 20) and 'Poll Rate Large' (set to 60). The 'Event Module' section has a checkbox labeled 'Enable/Disable the Solace Event Module'.

The **DATA COLLECTION** tab has the following fields:

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, SolAppliances and SolEnvironmentSensors.

DATA STORAGE

Use the **DATA STORAGE** tab to enable/disable the collection of historical data, modify the default settings for how data is collected and aggregated, as well as manage (or reduce) the amount of data stored.

There are several options for managing the amount of history data you store in your cache tables. The data for each metric is stored in a specific cache and, when the data is not updated in a given period of time, that data is either marked as *expired* or, if it has been an extended period of time, it is deleted from the cache altogether. By default, metric data is set to expired when the data in the cache has not been updated within **45** seconds. Also, by default, if the data is not updated in the cache within **3600** seconds, it is removed from the cache.

There are three main approaches to reducing the amount of storage in your cache tables. You can:

- **Set size limits:** So that when the number of rows in the cache table exceeds the limit specified here, the oldest row of data in the cache table is deleted. Set these values in the **History Rows** field.
- **Set time limits:** So that when the amount of time to wait for a row in the cache table to receive a data update exceeds the limit specified here, the data is marked expired or deleted. Set these values in the **Expire Time** and **Delete Time** fields.
- **Compact data:** Data compaction allows you to create rules that reduce the amount of stored history data to a reasonably sized sample of your data to prevent overloading your database. The rules include a schedule for automatically reducing the amount of data. Set these values in the **Condense Interval**, **Condense Raw Time** and **Compaction Rules** fields.

The screenshot displays the 'Solace Message Router' configuration page in the RTView Configuration Application. The interface is divided into a left sidebar and a main content area. The sidebar contains a 'Server Configuration' section with links to General, Databases, Alerts, Data Server, Historian, and Display Server. Below this is a 'Solution Package Configuration' section with a link to 'Solace Message Router'. The main content area has three tabs: 'CONNECTIONS', 'DATA COLLECTION', and 'DATA STORAGE'. The 'DATA STORAGE' tab is active, showing configuration options for history storage. The 'Size' section includes a 'History Rows' field set to 50000. The 'Compaction' section includes 'Condense Interval' (60), 'Condense Raw Time' (1200), and 'Compaction Rules' (1h - ;1d 5m ;2w 15m). The 'Duration' section includes 'Expire Time' (120), 'Delete Time' (3600), and 'Delete Time for Clients' (600). The 'History Storage' section lists various metrics with toggle switches and a 'Default' label: Message Routers (checked), Bridge Stats (unchecked), Client Stats (unchecked), CSPF Neighbors (checked), Endpoint Stats (unchecked), Endpoints (unchecked), Interface (unchecked), Message Spools (unchecked), Syslog Events (unchecked), and VPNs (checked). A 'History Table Name Prefix' field is at the bottom.

Solace Message Router

HOME SAVE

Server Configuration

- General
- Databases
- Alerts
- Data Server
- Historian
- Display Server

Solution Package Configuration

- Solace Message Router

CONNECTIONS DATA COLLECTION DATA STORAGE

Size
Set the number of history rows to keep in memory.

History Rows
50000

Compaction

Condense Interval	Condense Raw Time	Compaction Rules
60	1200	1h - ;1d 5m ;2w 15m

Duration
Set the number of seconds between data updates before metrics are expired or deleted.

Expire Time	Delete Time	Delete Time for Clients
120	3600	600

History Storage
Select metrics the Historian will store in the history database. Metrics that are not listed do not support storing history.

Metric	Status	Default
Message Routers	Checked	Default
Bridge Stats	Unchecked	Default
Client Stats	Unchecked	Default
CSPF Neighbors	Checked	Default
Endpoint Stats	Unchecked	Default
Endpoints	Unchecked	Default
Interface	Unchecked	Default
Message Spools	Unchecked	Default
Syslog Events	Unchecked	Default
VPNs	Checked	Default

History Table Name Prefix

The **DATA STORAGE** tab has the following fields:

- **Size** - Enter the maximum number of table rows to keep in memory:
 - History Rows** - The maximum number of rows to store in the History table. If set to **0**, no History table is created. The default setting is **50,000**. Caches impacted by this field are SolVpns, SolClientStats, SolAppliances, SolEndpoints, SolCspfNeighbors, SolBridgeStats, SolApplianceInterfaces, SolApplianceMessageSpool, SolEndpointStats and SolAppliancesQuality.
- **Compaction** - Define scheduled rules that reduce the amount of stored history data to a reasonably sized sample of your data to prevent overloading your database.
 - Condense Interval** - The time interval at which the cache history is condensed. The default is **60** seconds. Caches impacted by this field are SolVpns, SolClientStats, SolAppliances, SolEndpoints, SolCspfNeighbors, SolBridgeStats, SolApplianceInterfaces, SolApplianceMessageSpool and SolEndpointStatscaches.
 - Condense Raw Time** - The time span of raw data kept in the cache history table. The default is **1200** seconds. Caches impacted by this field are 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). Caches impacted by this field are SolVpns, SolClientStats, SolAppliances, SolEndpoints, SolCspfNeighbors, SolBridgeStats, SolApplianceInterfaces, SolApplianceMessageSpool and SolEndpointStats.
- **Duration** - Enter the amount of time between data updates before data is expired or deleted:
 - Expire Time** - The number of seconds to wait for a data update before cached history data is shown as **Expired** in displays. The caches impacted by this field are SolVpns, SolBridges, SolClients, SolClientStats, SolAppliances, SolEndpoints, SolCspfNeighbors, SolBridgeStats, SolApplianceInterfaces, SolApplianceMessageSpool, SolEndpointStats, SolEnvironmentSensors and SolAppliancesQuality.
 - Delete Time** - The number of seconds to wait for a data update before cached history data is removed from displays. The caches impacted by this field are SolVpns, SolBridges, SolEndpoints, SolBridgeStats, SolEndpointStats and SolEnvironmentSensors.
 - Delete Time for Clients** - The number of seconds to wait for a response from clients before removing client from displays. The caches impacted by this field are SolClients and SolClientStats.
- **History Storage** - Toggle to Enable/Disable the types of data you want the Historian to store in the history database for Solace. Each data type has a corresponding cache. Data types and the corresponding caches are:
 - Brokers:** Cache is SolAppliances.
 - Bridge Stats:** Cache is SolBridgeStats.
 - Client Stats:** Cache is SolClientStats.
 - CSPF Neighbors:** Cache is SolCspfNeighbors.
 - Endpoint Stats:** Cache is SolEndpointStats.
 - Endpoints:** Cache is SolEndpoints.

Interface: Cache is SolApplianceInterfaces.


Message Spools: Cache is SolApplianceMessageSpool.

Syslog Events: Cache is SyslogEvents

VPNs: Cache is SolVpns.

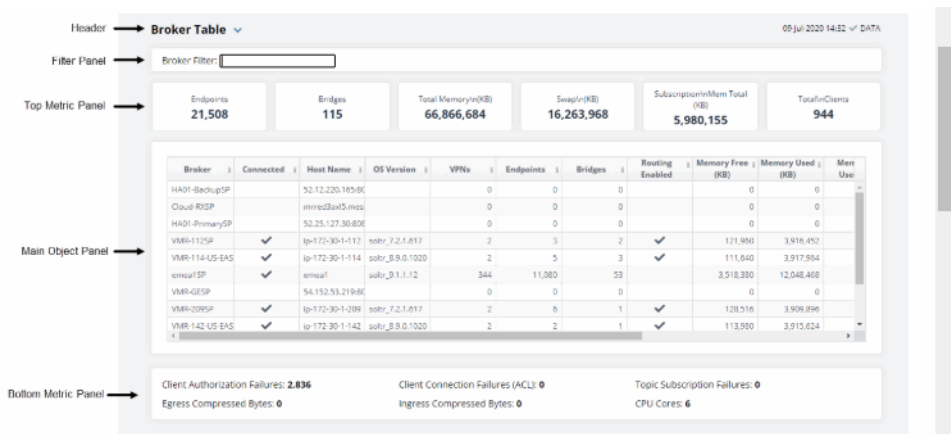
- **History Table Name Prefix** - Enter a prefix to prepend to the history data table names for these metrics.

Custom Display Designer

Use the custom display designer to quickly design and deploy custom displays in your monitor. You open the custom display designer by clicking the  button in the Configuration Application **HOME** page.

The panel on the left shows a list of your custom displays and the main panel allows you to edit those displays. To get started, click <New> in the display panel to start a new display. Once you have finished editing your display, click the Save button at the top of the display list. Navigate to the monitor and you will see a **CUSTOM** tab that contains your new display. If you do not see the **CUSTOM** tab, or you do not see your display in the navigation tree of your **CUSTOM** tab, click refresh in the monitor. Back in the designer, click <New> to create another new display or select an existing display to continue editing it.

The main panel allows you to edit your display. It is broken into 5 sections which correspond to the available sections in your custom display:



The screenshot shows the Custom Display Designer interface with the following components:

- Header:** Broker Table, 09 Jul 2020 14:02, DATA
- Filter Panel:** Broker Filter
- Top Metric Panel:**
 - Endpoints: 21,508
 - Bridges: 115
 - Total Memory(KB): 66,866,684
 - Swap(KB): 16,263,968
 - SubscriptionMem Total (KB): 5,980,155
 - Total Clients: 944
- Main Object Panel:**

Broker	Connected	Host Name	OS Version	VPNs	Endpoints	Bridges	Routing Enabled	Memory Free (KB)	Memory Used (KB)	Mem Use
HA01-BackupSP		52.12.220.165:06		0	0	0		0	0	
Cloud RISP		emred3ad5.mcs		0	0	0		0	0	
HA01-PrimarySP		52.25.127.30:06		0	0	0		0	0	
VMR-112SP	✓	ip-172-30-1-112	solr_2.2.1.617	2	3	2	✓	121,980	3,916,452	
VMR-114 US-EAS	✓	ip-172-30-1-114	solr_8.9.0.1020	2	5	3	✓	111,640	3,917,954	
emredSP	✓	emred1	solr_3.1.1.12	344	11,080	53		3,918,380	12,048,468	
VMR-GE3P	✓	54.152.53.219:06		0	0	0		0	0	
VMR-009SP	✓	ip-172-30-1-009	solr_2.2.1.617	2	6	1	✓	128,516	3,908,896	
VMR-142 US-EAS	✓	ip-172-30-1-142	solr_8.9.0.1020	2	2	1	✓	113,980	3,915,624	
- Bottom Metric Panel:**
 - Client Authorization Failures: 2,836
 - Client Connection Failures (ACL): 0
 - Topic Subscription Failures: 0
 - Egress Compressed Bytes: 0
 - Ingress Compressed Bytes: 0
 - CPU Cores: 6

Click the ? button in the top right corner of any section to display help text specific to that section. All Sections are optional and will only be included in the display if configured. The objects in the display can be configured to show data from any available solution package. However, if a display shows data from multiple solution packages, all of those solution packages must be hosted in the same RTViewDataServer.

Custom displays are stored in your installation under **projects/custom/displays**. Each display has 2 files: a **.txt** file and a corresponding **.html** file. To copy displays from one monitor to another, copy the whole displays directory into the **projects/custom** directory of the target monitor. When deployed with HA monitors, custom displays will be saved to the active monitor installation. You will need to copy them over to the backup monitor.

By default, all of the displays in the **projects/custom/displays** directory are shown in the **CUSTOM** tab of the monitor in a single level navigation tree. This is useful when developing custom displays since you can immediately see your changes in the **CUSTOM** tab as soon as you save them in the designer. Once you are ready to deploy, you can optionally create a custom navigation tree as follows:


In a text editor, create a file named `navtree.json` in the **projects/custom/displays** directory that contains a json array of objects where each object represents a node in the navigation tree and has the following fields:

- **label**: The label to use for the tree node.
- **display**: The name of the display to use for the tree node with the `.html` extension.
- **children**: Optional. An array of the nodes in the subtree expanded when this node is clicked.

For example, the following defines a 2 level navigation tree:

```
[
  {
    "label": "Summary A",
    "display": "a_top.html",
    "children": [
      {
        "label": "Detail A1",
        "display": "a1.html"
      },
      {
        "label": "Detail A2",
        "display": "a2.html"
      },
      {
        "label": "Detail A3",
        "display": "a3.html"
      }
    ]
  },
  {
    "label": "Summary B",
    "display": "b_top.html",
    "children": [
      {
        "label": "Detail B1",
        "display": "b1.html"
      },
      {
        "label": "Detail B2",
        "display": "b2.html"
      }
    ]
  }
]
```

That looks like this



You can also optionally hide or rename the **CUSTOM** tab as follows. In a text editor, create a file named `config.json` in the **projects/custom/displays** directory that contains the following json:

```
{
  "tabLabel": "Your Custom Label",
  "visible": true
}
```

Note that while you are developing displays, you will use the **CUSTOM** tab to see the displays you are editing. Therefore, you should wait to create a custom navigation tree until you are finished adding new displays. Or add new displays to your custom navigation tree as soon as you start working on them.

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**. Also on UNIX systems, it is a requirement that the installation directory path not contain spaces.

These instructions assume use of a BASH or a BASH-compliant shell.

Script Name	Description
my_alert_actions.bat/sh	Sample script to define actions for alerts. Location: The project directory. Format: my_alert_actions (Append .sh on UNIX)
rtv_setup.bat/sh	Initializes a command prompt or terminal window. Location: <installation directory>/bin This script must be executed in the directory in which it resides. Format: rtv_setup (Append .sh on UNIX)
rtvapm_init.bat/sh	Initializes a command window. Location: rtvapm This script must be executed in the directory in which it resides. Format: rtvapm_init (Append .sh on UNIX)

start_cmd.bat	<p>Starts an initialized Command Prompt window on Windows.</p> <p>Location:</p> <p><installation directory>/bin</p> <p>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.</p>
start_rtv.bat/sh	<p>Starts processes in an RTView configuration as specified in the rtvservers.dat configuration file.</p> <p>Location:</p> <p>rtvapm/common/bin</p> <p>This script must be executed in the project directory (the directory containing the rtvservers.dat file). This script requires rtvapm_init.bat/sh be executed first.</p> <p>An RTView configuration might include a Data Server or Display Server, an Historian and a Central Server Database. start_rtv only attempts to start processes it detects are not running. The action can be applied to all RTView configurations, a single RTView configuration or a single process in an RTView configuration.</p> <p>Before starting an RTView server, this script detects port conflicts caused by another server. If the conflict is caused by another RTView server, it returns a message identifying that server by its rtvapm. For example:</p> <pre>...start_rtv.bat: another dataserver running with JMX port 3268 under C:\rtview\RTViewDataServer\rtvapm</pre> <p>If the port conflict is caused by a non-RTView process, it returns a message similar to this, for example:</p> <pre>...start_rtv.bat: JMX port 3268 in use by PID 1234</pre> <p>In both cases the script includes this advice:</p> <p>Warning: server not started, port conflict</p> <hr/> <p>To avoid port conflicts, run your start script with the -portprefix: command line argument to change the first two (2) digits of all your server ports.</p> <p>To persist these port changes, change the port prefix in the RTView Configuration Application or use the -saveportprefix command line argument.</p> <p>Additional arguments can be included on the command line in which case they are passed to every server specified by the command.</p> <p>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.</p> <p>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.</p>

—console (or **—c**) - Start the processes with a command window (which is useful for testing).

When used without arguments, this script returns usage information and a list of available configurations. For example, **start_rtv** returns:

Usage: **start_rtv config or 'all' [server or 'all'] [args...]**

Available configs:

```

    default
        dataserver
        historian
        displayserver
        database
    sender
        dataserver

```

all

Starts all RTView configurations that are specified in the **rtvservers.dat** file.

all applies the action to all RTView configurations specified in the **rtvservers.dat** file (and corresponding servers or clients specified in each configuration). **Note:** When multiple configurations are specified in the **rtvservers.dat** file and they have different project settings directory locations, the **all** argument processes all the configurations. However, if the configurations have the same project settings directory locations, the **all** argument processes only the first configuration as the others are considered alternative configurations.

Example:

start_rtv all
(Append **.sh** on UNIX)

[Configuration Name]

Starts a single RTView configuration specified in the **rtvservers.dat** file:

start_rtv [Configuration Name]
(Append **.sh** on UNIX)

Configuration Name is the RTView configuration name specified in the **rtvservers.dat** file. The action applies to all servers or clients specified in the configuration.

Example:

start_rtv web_deployment
(Append **.sh** on UNIX)

[Server Name]

Starts a single process in an RTView configuration specified in the **rtvservers.dat** file:

start_rtv [Configuration Name] [Server Name]
(Append **.sh** on UNIX)

Server Name is the name of a server or client member in the configuration. For example, **dataserver**, **displayserver**, **historian** and **database**. The action applies only to that server or client in the configuration.

Example:

start_rtv web_deployment dataserver
(Append **.sh** on UNIX)

Use With Secured JMX Ports

This script works with RTView servers whose JMX ports are secured with either a username and password, or with SSL. You provide the scripts with the necessary credential information and the scripts manage authentication with the server. There are two ways that you can provide credential information to the scripts: via command-line arguments and via properties placed in any property file that is used by the server.

Securing with username and password

- To secure with a username and password via command-line, use the arguments as follows:

-jmxuser:...

-jmxpass:...

- To secure with a username and password in a property file, use the properties as follows:

sl.rtvview.jmxremote.username=...

sl.rtvview.jmxremote.password=....

Securing with SSL

To secure with SSL, you provide the client KeyStore and TrustStore locations and their corresponding passwords.

- To secure with SSL via command-line, use the arguments as follows:

-sslkeystore:...

-sslkeystorepass:...

-ssltruststore:...

-ssltruststorepass:...

- To secure with SSL in a property file, use the properties as follows:

sl.rtvview.ssl.client.keyStore=...

sl.rtvview.ssl.client.keyStorePassword=...

sl.rtvview.ssl.client.trustStore=...

sl.rtvview.ssl.client.trustStorePassword=....

Password Encryption

To encrypt the passwords in your properties files, use the command-line tool "encode_string", for example:

encode_string encoder2 password

This will give you an encrypted value for "password" that you can use in your properties.

start_server.bat/sh

Starts the RTView DataServer.

Location:

<installation directory>

This script must be executed in the directory in which it resides. You can also execute the script by double-clicking in an Explorer window.

Format:

start_server

(Append .sh on UNIX)

start_servers.bat/sh	<p>Starts the RTViewCentral servers.</p> <p>Location: <installation directory>/bin</p> <p>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.</p> <p>Format: start_servers (Append .sh on UNIX)</p>
start_tomcat.bat/sh	<p>Starts Apache Tomcat.</p> <p>Location: <installation directory>/bin</p> <p>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.</p> <p>Format: start_tomcat (Append .sh on UNIX)</p>
status_collector.bat/sh	<p>Returns the status of RTView DataCollector.</p> <p>Location: <installation directory></p> <p>This script must be executed in the project directory (the directory containing the rtvservers.dat file).</p> <p>Format: status_collector (Append .sh on UNIX)</p>
status_rtv.bat/sh	<p>Returns the status of all RTView configurations that are specified in the rtvservers.dat configuration file.</p> <p>Location: rtvapm/common/bin</p> <p>This script must be executed in the project directory (the directory containing the rtvservers.dat file). This script requires rtvapm_init.bat/sh be executed first.</p> <p>This action uses defined JMX ports. An RTView configuration might include a Data Server, a Display Server or Viewer, an Historian and a Central Server Database. status_rtv only attempts to start processes it detects are not running. The action can be applied to all RTView configurations, a single RTView configuration or a single process in an RTView configuration.</p> <p>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.</p> <p>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.</p>

all

Returns the status of all RTView configurations specified in the **rtvservers.dat** file. **Note:** When multiple configurations are specified in the **rtvservers.dat** file and they have different project settings directory locations, the **all** argument processes all the configurations. However, if the configurations have the same project settings directory locations, the **all** argument processes only the first configuration as the others are considered alternative configurations.

Example:

status_rtv all
(Append **.sh** on UNIX)

[Configuration Name]

Returns the status of a single RTView configuration specified in the **rtvservers.dat** file:

status_rtv [Configuration Name]
(Append **.sh** on UNIX)

Configuration Name is the RTView configuration name specified in the **rtvservers.dat** file. The action applies to all servers or clients specified in the configuration.

Example:

status_rtv web_deployment
(Append **.sh** on UNIX)

[Server Name]

Returns the status of a single process in an RTView configuration specified in the **rtvservers.dat** file:

status_rtv [Configuration Name] [Server Name]
(Append **.sh** on UNIX)

Server Name is the name of a server or client member in the configuration. For example, **dataserver**, **displayserver**, **historian** and **database**. The action applies only to that server or client in the configuration.

Example:

status_rtv web_deployment dataserver
(Append **.sh** on UNIX)

Use With Secured JMX Ports

This script works with RTView servers whose JMX ports are secured with either a username and password, or with SSL. You provide the scripts with the necessary credential information and the scripts manage authentication with the server. There are two ways that you can provide credential information to the scripts: via command-line arguments and via properties placed in any property file that is used by the server.

Securing with username and password

- To secure with a username and password via command-line, use the arguments as follows:

-jmxuser:...

-jmxpass:...

- To secure with a username and password in a property file, use the properties as follows:

sl.rtvview.jmxremote.username=...

sl.rtvview.jmxremote.password=....

Securing with SSL

To secure with SSL, you provide the client KeyStore and TrustStore locations and their corresponding passwords.

- To secure with SSL via command-line, use the arguments as follows:

-sslkeystore:...

-sslkeystorepass:...

-ssltruststore:...

-ssltruststorepass:...

- To secure with SSL in a property file, use the properties as follows:

sl.rtvview.ssl.client.keyStore=...

sl.rtvview.ssl.client.keyStorePassword=...

sl.rtvview.ssl.client.trustStore=...

sl.rtvview.ssl.client.trustStorePassword=....

Password Encryption

To encrypt the passwords in your properties files, use the command-line tool "encode_string", for example:

encode_string encoder2 password

This will give you an encrypted value for "password" that you can use in your properties.

status_server.bat/sh

Returns the status of the RTView DataServer.

Location:

<installation directory>

This script must be executed in the project directory (the directory containing the **rtvservers.dat** file).

Format:

status_server

(Append **.sh** on UNIX)

status_servers.bat/sh	<p>Returns the status of the RTViewCentral servers (as well as the Solace PubSub+ Monitor in RTViewSolaceMonitor).</p> <p>Location: <installation directory>/bin</p> <p>This script must be executed in the project directory (the directory containing the rtvservers.dat file).</p> <p>Format: status_servers (Append .sh on UNIX)</p>
stop_collector.bat/sh	<p>Stops the RTView DataCollector.</p> <p>Location: <installation directory></p> <p>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.</p> <p>Format: stop_collector (Append .sh on UNIX)</p>
stop_rtv.bat/sh	<p>Stops processes in an RTView configuration as specified in the rtvservers.dat configuration file.</p> <p>Location: rtvapm/common/bin</p> <p>This script must be executed in the project directory (the directory containing the rtvservers.dat file). This script requires rtvapm_init.bat/sh be executed first.</p> <p>This action uses defined JMX ports. An RTView configuration might include a Data Server or a Display Server, an Historian and a Central Server Database. stop_rtv only attempts to start processes it detects are not running. The action can be applied to all RTView configurations, a single RTView configuration or a single process in an RTView configuration.</p> <p>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.</p> <p>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.</p> <p>Location: project directory</p> <p>This script must be executed in the project directory (the directory containing the rtvservers.dat file).</p>

all

Stops all RTView configurations that are specified in the **rtvservers.dat** file. **all** applies the action to all RTView configurations specified in the **rtvservers.dat** file (and corresponding servers or clients specified in each configuration).

Note: When multiple configurations are specified in the **rtvservers.dat** file and they have different project settings directory locations, the **all** argument processes all the configurations. However, if the configurations have the same project settings directory locations, the **all** argument processes only the first configuration as the others are considered alternative configurations.

Example:

stop_rtv all
(Append **.sh** on UNIX)

[Configuration Name]

Stops a single RTView configuration specified in the **rtvservers.dat** file:

stop_rtv [Configuration Name]
(Append **.sh** on UNIX)

Configuration Name is the RTView configuration name specified in the **rtvservers.dat** file. The action applies to all servers or clients specified in the configuration.

Example:

stop_rtv web_deployment
(Append **.sh** on UNIX)

[Server Name]

Stops a single process in an RTView configuration specified in the **rtvservers.dat** file:

stop_rtv [Configuration Name] [Server Name]
(Append **.sh** on UNIX)

Server Name is the name of a server or client member in the configuration. For example, **dataserver**, **displayserver**, **historian** and **database**. The action applies only to that server or client in the configuration.

Example:

stop_rtv web_deployment dataserver
(Append **.sh** on UNIX)

Use With Secured JMX Ports

This script works with RTView servers whose JMX ports are secured with either a username and password, or with SSL. You provide the scripts with the necessary credential information and the scripts manage authentication with the server. There are two ways that you can provide credential information to the scripts: via command-line arguments and via properties placed in any property file that is used by the server.

Securing with username and password

- To secure with a username and password via command-line, use the arguments as follows:

-jmxuser:...

-jmxpass:...

- To secure with a username and password in a property file, use the properties as follows:

sl.rtvview.jmxremote.username=...

sl.rtvview.jmxremote.password=....

Securing with SSL

To secure with SSL, you provide the client KeyStore and TrustStore locations and their corresponding passwords.

- To secure with SSL via command-line, use the arguments as follows:

-sslkeystore:...

-sslkeystorepass:...

-ssltruststore:...

-ssltruststorepass:...

- To secure with SSL in a property file, use the properties as follows:

sl.rtvview.ssl.client.keyStore=...

sl.rtvview.ssl.client.keyStorePassword=...

sl.rtvview.ssl.client.trustStore=...

sl.rtvview.ssl.client.trustStorePassword=....

Password Encryption

To encrypt the passwords in your properties files, use the command-line tool "encode_string", for example:

encode_string encoder2 password

This will give you an encrypted value for "password" that you can use in your properties.

stop_server.bat/sh

Stops the RTView DataServer.

Location:

<installation directory>

This script must be executed in the directory in which it resides.

Format:

stop_server

(Append .sh on UNIX)

stop_servers.bat/sh	<p>Stops the RTViewCentral servers.</p> <p>Location: <installation directory>/bin</p> <p>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.</p> <p>Format: stop_servers (Append .sh on UNIX)</p>
stop_tomcat.bat/sh	<p>Stops Apache Tomcat.</p> <p>Location: <installation directory>/bin</p> <p>This script must be executed in the directory in which it resides.</p> <p>Format: start_tomcat (Append .sh on UNIX)</p>
update_wars.bat/sh	<p>Creates/updates the primary Monitor servlets.</p> <p>Location: <installation directory>/projects/rtview-server</p> <p>This script must be executed in the directory in which it resides. This script requires rtvapm_init.bat/sh be executed first.</p> <p>Format: update_wars.sh [appname [host [portprefix]]]</p> <p>For example: update_wars.sh my-appname my-hostname 99</p> <p>The name, host, and portprefix are declared in variables at the top of the script for easy editing, and can be passed into the scripts on the command-line.</p> <p>-secure Use the "-secure" argument to update the rtvquery war with security enabled.</p> <p>You can use ? or help to get a usage message. For example: update_wars.sh help</p> <p>You can edit other variables at the top of the scripts to set properties for high-availability (HA).</p> <p>Set HA_HOST to the hostname of the backup data server.</p> <p>Set HA_DISPLAYHOST to the hostname of the backup display server.</p> <p>Set HA_FAILBACK to true to automatically reconnect to the primary display server.</p>

validate_install.bat/sh	<p>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.</p> <p>Location:</p> <p><installation directory>/bin</p> <p>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.</p> <p>The script returns the following information (where <RTViewInstallation> is your RTView installation):</p> <ul style="list-style-type: none"> • In Windows <p>Validating installation in /opt/rtview/<RTViewInstallation> ... Java installation correct. ... rtvapm installation correct.</p> <ul style="list-style-type: none"> • In UNIX <p>Validating installation in /opt/rtview/<RTViewInstallation> ... Java installation correct. ... rtvapm installation correct. ... file permissions correct. ... file formats correct.</p>
--------------------------------	---

rtvservers.dat

This section describes the **rtvservers.dat** configuration file which is used to manage your RTView deployment and processes. This section includes:

- ["Single Configuration File"](#)
- ["Multiple Configuration File"](#)

The **rtvservers.dat** text file contains one or more RTView configurations. An RTView configuration is a group of servers that should be started together. For example, the configuration might include any of the following: a Data Server, Historian, HSQLDB database, and a Display Server (for a Web Deployment). The **rtvservers.dat** file is used when the following scripts are executed:

- [start_rtv](#) Starts RTView processes specified in the **rtvservers.dat** file.
- [stop_rtv](#) Stops the RTView processes specified in the **rtvservers.dat** file.
- [status_rtv](#) Returns status information for RTView processes specified in the **rtvservers.dat** file.

Single Configuration File

The following **rtvservers.dat** file, located in your project directory, contains a single RTView configuration, named **default**.

```
default . dataserver rundata
```

```
default . historian runhist -ds
default . displayserver rundisp -ds
default . database rundb
```

Note: The last line in the **rtvservers.dat** file must end with a new line, or be followed by a blank line.

In this example, to start the **default** configuration type: **start_rtv default** or **start_rtv all**. To start a single server in the configuration, type **start_rtv <Configuration Name> <Server Name>**. For example: **start_rtv default displayserver**.

Each line has the following format consisting of four fields:

<Configuration Name> <Project Settings Directory Location> <Property Filter Identifying the Server> <Command>

<Configuration Name>	The name of the RTView configuration (default in this example).
<Project Settings Directory Location>	The RTView project settings directory location, relative to the location of the rtvservers.dat file (., the current directory, in this example).
<Property Filter Identifying the Server>	The property filter that identifies the server, which is the property filter under which the server's JMX port is defined. By default, this is the server name, such as dataserver , displayserver and historian .
<Command>	The script used to start the process. Valid values are: <ul style="list-style-type: none"> • 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 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 [“Alerts for Solution Package for Solace”](#) and [“Alerts for RTView Manager”](#).

Alerts for Solution Package for Solace

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
SolBrokerEgressByteRateHigh The egress rate (bytes/sec) for the message broker is excessive. Index Type: PerAppliance	70	85	30	FALSE
SolBrokerIngressByteRateHigh The ingress rate (bytes/sec) for the message broker is excessive. Index Type: PerAppliance	70	85	30	FALSE
SolBrokerEgressMsgRateHigh The message egress rate (the number of egress messages per second) for the message broker is excessive. Index Type: PerAppliance	70	85	30	FALSE

SolBrokerIngressMsgRateHigh The message ingress rate (the number of ingress messages per second) for the message broker is excessive. Index Type: PerAppliance	70	85	30	FALSE
SolBrokerNoQueueFound This is an Event Alert. Event Alerts do not have duration or threshold settings. A single alarm alert when there are discarded queues in the broker. (Delta of discard-queue-not-found is non-zero). Note: This alert cannot be executed for Cloud Brokers. This request XML is a system level request which means that Cloud login credentials do not have permission to execute it. Index Type: PerBroker				FALSE
SolBrokerNoSubscriptionMatch This is an Event Alert. Event Alerts do not have duration or threshold settings. A single alarm alert when there are no current subscription matches (Delta of no-subscription-match is non-zero). Note: This alert cannot be executed for Cloud Brokers. This request XML is a system level request which means that Cloud login credentials do not have permission to execute it. Index Type: PerBroker				FALSE
SolBrokerNoValidDestination This is an Event Alert. Event Alerts do not have duration or threshold settings. A single alarm alert when invalid destinations exist in the broker. (Delta of discard-nodest is non-zero). Note: This alert cannot be executed for Cloud Brokers. This request XML is a system level request which means that Cloud login credentials do not have permission to execute it. Index Type: PerBroker				FALSE

SolBrokerRedundancyDown & SolBrokerRedundancyActivityStatusChanged

These alerts only pertain to brokers that are configured for redundancy.

These alerts execute when a redundancy misconfiguration is detected. Brokers qualify as being configured for redundancy if the Monitor either detects an associated mate broker name or a broker is explicitly configured for redundancy.

To verify whether the Monitor has detected all brokers configured for redundancy, go to the **Admin>Cache Table** display and select the **_SolBrokerRedundancy** cache. Verify that all brokers that are configured for redundancy have the **IsHABroker** flag checked. If the **IsHABroker** flag is NOT checked, use the RTView Configuration Application to configure the brokers for redundancy.

The **SolBrokerRedundancyDown** alert verifies that redundancy is configured properly by checking whether **Redundancy Mode**, **Redundancy Status** and **Configuration Status** are valid. That is, the **Redundancy Mode** is either **Active/Active** or **Active/Standby**, the **Redundancy Status** is **Up** and the **Redundancy Configuration Status** is **Enabled**. If any of these conditions are not met, then a warning alert will be raised with the following alert text: "<hostname> is not properly configured for redundancy or redundancy is down. Redundancy Status: <a> Configuration Status: ", where <hostname> is the hostname of the offending broker and <a> and are the current Redundancy Status and Configuration Status of the broker respectively.

The **SolBrokerRedundancyActivityStatusChanged** alert checks whether the previous state of the **Active-Standby Role**, the **Activity Status** of the Primary Virtual Router and the **Activity Status** of the Backup Virtual Router is different from the current state. If they are different, that implies a change in the state of the redundancy status occurred and a warning alert will be triggered. As soon as the previous and the current redundancy state is stabilized, the warning alert automatically clears, indicating in the alert text the current and previous states being detected. The warning alert contains the following text: "<hostname> has changed its redundancy activity state. There might be untracked intermediate states from the ones that have been detected. Current state: <A> Previous state: ", where <A> and are the concatenation of active-standby-role, primary-status-activity, and backup-status-activity separated by the character "-" for current and previous states.

Best Practices & Troubleshooting

It's possible to have multiple **SolBrokerRedundancyActivityStatusChanged** warning alerts when failing over if intermediate states have been collected. For instance, if the changes from Local Active to Local Inactive to Shutdown are detected, then two **SolBrokerRedundancyActivityStatusChanged** warning alerts will be executed in this broker and will have two warnings from one broker and one from the other broker if the intermediate state on the second broker was not gathered due to polling interval being longer than the time the broker changes its redundancy state. If you only want one warning alert per broker per failover operation, the recommended action is to increase the duration of the alert. This value will vary depending on data collection latency and is system dependent. On the other hand, if you need to keep track of all intermediate states of the failover operation, then you should decrease the polling interval for the show redundancy detail poller. This is not recommended as might overflow the data collector with requests that cannot be successfully completed or preventing sending other monitoring data regarding other aspects of the broker due to the existence of requests too-often repeated.

Due to **SolBrokerRedundancyActivityStatusChanged** warning alert being a transient alert which will be automatically cleared when the redundancy status is stabilized, enabling both alerts is recommended as **SolBrokerRedundancyDown** can stay uncleared if manual intervention for fixing redundancy misconfiguration or non-functioning is required.

By default, these alerts are disabled.

SolClientInboundByteRateHigh The number of inbound bytes per second for the client has reached its maximum. Index Type: PerClient	8000000	10000000	30	FALSE
SolClientInboundMsgRateHigh The number of inbound 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
SolCspfNeighborDown State is not "OK" for one or more CSPF neighbors. Index Type: PerNeighbor	1	NaN	30	FALSE
SolEndpointNoBridgeClient This is an Event Alert. Event Alerts do not have duration or threshold settings. A single alarm alert when there are no binds for the Solace Endpoint exist (bind-count is zero). Index Type: PerEndpoint	NaN	NaN	NaN	FALSE
SolEndpointNoBridgeTopic This is an Event Alert. Event Alerts do not have duration or threshold settings. A single alarm alert when there are no topics subscribed to the Queue (topic-subscription-count is zero). Index Type: PerEndpoint	NaN	NaN	NaN	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 broker memory for storing spooled messages. (Threshold units are megabytes.) Index Type: PerEndpoint	40	50	30	FALSE
SolEventModuleBrokerAlert This is an Event Alert. Event Alerts do not have duration or threshold settings. If the Solace Event Module is properly configured and running and this alert is enabled, all Syslog Events that are selected as alerts from the Message Brokers that were enabled for being monitored with Syslog will trigger this type of alert from the SYSTEM scope. Alerts of this type refer to Syslog events that can be clearable and non-clearable of SYSTEM scope. Therefore this alert can be clearable and non-clearable, depending on the event that triggered its execution.				FALSE
SolEventModuleClientAlert This is an Event Alert. Event Alerts do not have duration or threshold settings. If the Solace Event Module is properly configured and running and this alert is enabled, all Syslog Events that are selected as alerts from the Message Brokers that were enabled for being monitored with Syslog will trigger this type of alert from the CLIENT scope. Alerts of this type refer to Syslog events that can be clearable and non-clearable of CLIENT scope. Therefore this alert can be clearable and non-clearable, depending on the event that triggered its execution.				FALSE

SolEventModuleVpnAlert This is an Event Alert. Event Alerts do not have duration or threshold settings. If the Solace Event Module is properly configured and running and this alert is enabled, all Syslog Events that are selected as alerts from the Message Brokers that were enabled for being monitored with Syslog will trigger this type of alert from the VPN scope. Alerts of this type refer to Syslog events that can be clearable and non-clearable of VPN scope. Therefore this alert can be clearable and non-clearable, depending on the event that triggered its execution.				FALSE
SolGuaranteedMsgingHbaLinkDown For Guaranteed Messaging only, the Operational State for each HBA Fibre-Channel should be Online (e.g., not Linkdown). Index Type: PerHbaLink	NaN	0	30	FALSE
SolGuaranteedMsgingMatePortDown This alert executes when the broker is configured for Guaranteed Messaging and the ADB Link to Mate is down. Note that this alert does not clear until a failover occurs. Index Type: PerADB	NaN	0	30	FALSE
SolGuaranteedMsgingNoMsgSpoolAdActive This alert applies to a pair of brokers that are configured for redundancy as an HA pair. A single alert executes when neither broker in the HA pair has a message spool operational status of AD-Active . Index Type: PerMsgRouter	NaN	0	30	FALSE
SolMsgBrokerExpired This is an Event Alert. Event Alerts do not have duration or threshold settings. The collection of monitoring data has stopped abruptly. Use this alert in conjunction with the SolMsgBrokerNotConnected alert, or instead of it, if you don't need to be notified about a lack of connection when the Monitor starts up.	NaN	NaN	NaN	FALSE
SolMsgRouterActiveDiskUtilHigh The utilization of the active disk partition for the message broker is excessive. Index Type: PerAppliance	70	85	30	FALSE
SolMsgRouterConnectionUtilHigh The connection utilization for the message broker (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

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 broker has reached its max threshold. Index Type: PerAppliance	400000	500000	30	FALSE
SolMsgRouterInboundMsgRateHigh The number of inbound messages per second for the message broker 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 broker 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 broker 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 broker 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 broker has reached its max threshold. Index Type: PerAppliance	400000	500000	30	FALSE
SolMsgRouterOutboundMsgRateHigh The number of outbound messages per second for the message broker has reached its max threshold. Index Type: PerAppliance	400000	500000	30	FALSE
SolMsgRouterPendingMsgsHigh The total number of pending messages for this message broker 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 percentage of spool spaces used for storing spooled messages is excessive. Index Type: PerAppliance	70	85	30	FALSE
SolMsgRouterStandbyDiskUtilHigh The utilization of the standby disk partition for the message broker is excessive. Index Type: PerAppliance	70	85	30	FALSE
SolMsgRouterSubscriptionUtilHigh The Topic Subscriptions on Queue utilization (current number of subscriptions divided by max allowed) for the message router is excessive. Index Type: PerAppliance	70	85	30	FALSE
SolMsgRouterSwapUsedHigh This alert strictly applies to software brokers (it does not execute for hardware brokers). The amount of swap space used by the message broker operating system is excessive. 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, 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 broker is excessive. The metrics are: (transacted-sessions-used/ max-transacted-sessions)*100 Index Type: PerMsgRouter	70	85	30	FALSE
SolMsgRouterTranSessionResUtilHigh The transacted session resource utilization for the message broker 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
SolSparseMessageSpoolFile This is a Limits Alert that issues a Warning alert and is enabled by default. Important: Do not modify thresholds for this alert as are set by Solace development. A single warning alert (Severity 1) executes when the active-disk-partition-usage > 30.0 AND disk-usage-mb/current-disk-usage >= 3.0. This alert is defined to determine when there is a Sparse Message Spool File Condition. When disk space usage is several multiples of persistent store usage, then there is likely a large number of message spool files residing on the disk where each file contains few messages. This is referred to as a sparse message spool file condition, and requires urgent attention to mitigate and avoid the disk reaching capacity. For further information, refer to Solace documentation for diagnosing the sparse message spool file condition.				TRUE

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

Alerts for RTView Manager

RTView Manager comes with the following alert types for RTView Servers (Data Servers, Display Servers and Historian Servers):

JvmCpuPercentHigh	<p>Executes a single warning alert and a single alarm alert if the percent of JVM CPU used exceeds the specified threshold.</p> <p>Index Type: Per JVM</p> <p>Metric: CpuPercent</p>
JvmGcDutyCycleHigh	<p>Executes a single warning alert and a single alarm alert if the garbage collector duty cycle exceeds the specified threshold.</p> <p>Index Type: Per GC Source</p> <p>Metric: DutyCycle</p>
JvmMemoryUsedAfterGCHigh	<p>Executes a single warning alert and a single alarm alert if the percent of memory used after garbage collection exceeds the specified threshold.</p> <p>Index Type: Per GC Source</p> <p>Metric: PctMemoryUsedAfterGC</p>
JvmMemoryUsedHigh	<p>Executes a single warning alert and a single alarm alert if the percent of memory used exceeds the specified threshold.</p> <p>Index Type(s): Per JVM</p> <p>Metric: MemoryUsedPercent</p>
JvmNotConnected	<p>Executes a single alert if the JVM is disconnected, indicating that it might have crashed.</p> <p>Index Type(s): Per JVM</p> <p>Metric: Connected</p>
JvmStaleData	<p>Executes a single alert if the data update wait time exceeds the specified duration threshold.</p> <p>Index Type(s): Per JVM</p> <p>Metric: Expired</p>
JvmThreadCountHigh	<p>Executes a single warning alert and a single alarm alert if the number of threads exceeds the specified threshold.</p> <p>Index Type(s): Per JVM</p> <p>Metric: ThreadCount</p>
TomcatAccessRateHigh	<p>Executes a single warning alert and a single alarm alert if the number of accesses per second exceeds the specified threshold.</p> <p>Index Type(s): Per Server</p> <p>Metric: RateaccessCount</p>
TomcatActiveSessionsHigh	<p>Executes a single warning alert and a single alarm alert if the number of active sessions exceeds the specified threshold.</p> <p>Index Type(s): Per Server</p> <p>Metric: activeSessions</p>

TomcatAppAccessRateHigh

Executes a single warning alert and a single alarm alert if the number of accesses per second exceeds the specified threshold.

Index Type(s): Per Application

Metric: RateaccessCount

TomcatAppActiveSessionsHigh

Executes a single warning alert and a single alarm alert if the number of active sessions exceeds the specified threshold.

Index Type(s): Per Application

Metric: activeSessions

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

** Apache Jackson libraries are licensed under the Apache License Version 2.11.0

** JEval 0.9.4 is licensed under the Apache License Version 2.0

** Jetty 9.4.19 is licensed under the Apache License Version 2.0

Apache License 2.0

Version 2.0, January 2004

<http://www.apache.org/licenses/>

TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

1. Definitions.

"License" shall mean the terms and conditions for use, reproduction, and distribution as defined by Sections 1 through 9 of this document.

"Licensor" shall mean the copyright owner or entity authorized by the copyright owner that is granting the License.

"Legal Entity" shall mean the union of the acting entity and all other entities that control, are controlled by, or are under common control with that entity. For the purposes of this definition, "control" means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity.

"You" (or "Your") shall mean an individual or Legal Entity exercising permissions granted by this License.

"Source" form shall mean the preferred form for making modifications, including but not limited to software source code, documentation source, and configuration files.

"Object" form shall mean anyform resulting from mechanical transformation or translation of a Source form, including but not limited to compiled object code, generated documentation, and conversions to other media types.

"Work" shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in or attached to the work (an example is provided in the Appendix below)

"Derivative Works" shall mean any work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial revisions, annotations, elaborations, or other modifications represent, as a whole, an original work of authorship. For the purposes of this License, Derivative Works shall not include works that remain separable from, or merely link (or bind by name) to the interfaces of, the Work and Derivative Works thereof.

"Contribution" shall mean any work of authorship, including the original version of the Work and any modifications or additions to that Work or Derivative Works thereof, that is intentionally submitted to Licensor for inclusion in the Work by the copyright owner or by an individual or Legal Entity authorized to submit on behalf of the copyright owner. For the purposes of this definition, "submitted" means any form of electronic, verbal, or written communication sent to the Licensor or its representatives, including but not limited to communication on electronic mailing lists, source code control systems, and issue tracking systems that are managed by, or on behalf of, the Licensor for the purpose of discussing and improving the Work, but excluding communication that is conspicuously marked or otherwise designated in writing by the copyright owner as "Not a Contribution."

"Contributor" shall mean Licensor and any individual or Legal Entity on behalf of whom a Contribution has been received by Licensor and subsequently incorporated within the Work.

2. Grant of Copyright License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable copyright license to reproduce, prepare Derivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Source or Object form.

3. Grant of Patent License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable (except as stated in this section) patent license to make, have made, use, offer to sell, sell, import, and otherwise transfer the Work, where such license applies only to those patent claims licensable by such Contributor that are necessarily infringed by their Contribution(s) alone or by combination of their Contribution(s) with the Work to which such Contribution(s) was submitted. If You institute patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent infringement, then any patent licenses granted to You under this License for that Work shall terminate as of the date such litigation is filed.

4. Redistribution. You may reproduce and distribute copies of the Work or Derivative Works thereof in any medium, with or without modifications, and in Source or Object form, provided that You meet the following conditions:

- (a) You must give any other recipients of the Work or Derivative Works a copy of this License; and
- (b) You must cause any modified files to carry prominent notices stating that You changed the files; and
- (c) You must retain, in the Source form of any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form of the Work, excluding those notices that do not pertain to any part of the Derivative Works; and

(d) If the Work includes a "NOTICE" text file as part of its distribution, then any Derivative Works that You distribute must include a readable copy of the attribution notices contained within such NOTICE file, excluding those notices that do not pertain to any part of the Derivative Works, in at least one of the following places: within a NOTICE text file distributed as part of the Derivative Works; within the Source form or documentation, if provided along with the Derivative Works; or, within a display generated by the Derivative Works, if and wherever such third-party notices normally appear. The contents of the NOTICE file are for informational purposes only and do not modify the License. You may add Your own attribution notices within Derivative Works that You distribute, alongside or as an addendum to the NOTICE text from the Work, provided that such additional attribution notices cannot be construed as modifying the License.

You may add Your own copyright statement to Your modifications and may provide additional or different license terms and conditions for use, reproduction, or distribution of Your modifications for any such Derivative Works as a whole, provided Your use, reproduction, and distribution of the Work otherwise complies with the conditions stated in this License.

5. Submission of Contributions. Unless You explicitly state otherwise, any Contribution intentionally submitted for inclusion in the Work by You to the Licensor shall be under the terms and conditions of this License, without any additional terms or conditions. Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions.

6. Trademarks. This License does not grant permission to use the trade names, trademarks, service marks, or product names of the Licensor, except as required for reasonable and customary use in describing the origin of the Work and reproducing the content of the NOTICE file.

7. Disclaimer of Warranty. Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License.

8. Limitation of Liability. In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, including any direct, indirect, special, incidental, or consequential damages of any character arising as a result of this License or out of the use or inability to use the Work (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses), even if such Contributor has been advised of the possibility of such damages.

9. Accepting Warranty or Additional Liability. While redistributing the Work or Derivative Works thereof, You may choose to offer, and charge a fee for, acceptance of support, warranty, indemnity, or other liability obligations and/or rights consistent with this License. However, in accepting such obligations, You may act only on Your own behalf and on Your sole responsibility, not on behalf of any other Contributor, and only if You agree to indemnify, defend, and hold each Contributor harmless for any liability incurred by, or claims asserted against, such Contributor by reason of your accepting any such warranty or additional liability.

END OF TERMS AND CONDITIONS

APPENDIX: How to apply the Apache License to your work.

To apply the Apache License to your work, attach the following boilerplate notice, with the fields enclosed by brackets "[]" replaced with your own identifying information. (Don't include the brackets!) The text should be enclosed in the appropriate comment syntax for the file format. We also recommend that a file or class name and description of purpose be included on the same "printed page" as the copyright notice for easier identification within third-party archives.

Copyright [yyyy] [name of copyright owner]

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at:

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

=====

** TreeMap Algorithms v1.0 is used without modifications and licensed by MPL Version 1.1. The source for TreeMap Algorithms can be obtained from <http://www.cs.umd.edu/hcil/treemap/>

** iTextAsian 1.0 is licensed by MPL Version 1.1 and the source can be obtained from: <http://itextpdf.com/download.php>

MOZILLA PUBLIC LICENSE

Version 1.1

1. Definitions.

1.0.1. "Commercial Use" means distribution or otherwise making the Covered Code available to a third party.

1.1. "Contributor" means each entity that creates or contributes to the creation of Modifications.

1.2. "Contributor Version" means the combination of the Original Code, prior Modifications used by a Contributor, and the Modifications made by that particular Contributor.

1.3. "Covered Code" means the Original Code or Modifications or the combination of the Original Code and Modifications, in each case including portions thereof.

1.4. "Electronic Distribution Mechanism" means a mechanism generally accepted in the software development community for the electronic transfer of data.

1.5. "Executable" means Covered Code in any form other than Source Code.

1.6. "Initial Developer" means the individual or entity identified as the Initial Developer in the Source Code notice required by Exhibit A.

1.7. "Larger Work" means a work which combines Covered Code or portions thereof with code not governed by the terms of this License.

1.8. "License" means this document.

1.8.1. "Licensable" means having the right to grant, to the maximum extent possible, whether at the time of the initial grant or subsequently acquired, any and all of the rights conveyed herein.

1.9. "Modifications" means any addition to or deletion from the substance or structure of either the Original Code or any previous Modifications. When Covered Code is released as a series of files, a Modification is:

A. Any addition to or deletion from the contents of a file containing Original Code or previous Modifications.

B. Any new file that contains any part of the Original Code or previous Modifications.

1.10. "Original Code" means Source Code of computer software code which is described in the Source Code notice required by Exhibit A as Original Code, and which, at the time of its release under this License is not already Covered Code governed by this License.

1.10.1. "Patent Claims" means any patent claim(s), now owned or hereafter acquired, including without limitation, method, process, and apparatus claims, in any patent Licensable by grantor.

1.11. "Source Code" means the preferred form of the Covered Code for making modifications to it, including all modules it contains, plus any associated interface definition files, scripts used to control compilation and installation of an Executable, or source code differential comparisons against either the Original Code or another well known, available Covered Code of the Contributor's choice. The Source Code can be in a compressed or archival form, provided the appropriate decompression or de-archiving software is widely available for no charge.

1.12. "You" (or "Your") means an individual or a legal entity exercising rights under, and complying with all of the terms of, this License or a future version of this License issued under Section 6.1. For legal entities, "You" includes any entity which controls, is controlled by, or is under common control with You. For purposes of this definition, "control" means (a) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (b) ownership of more than fifty percent (50%) of the outstanding shares or beneficial ownership of such entity.

2. Source Code License.

2.1. The Initial Developer Grant.

The Initial Developer hereby grants You a world-wide, royalty-free, non-exclusive license, subject to third party intellectual property claims:

(a) under intellectual property rights (other than patent or trademark) Licensable by Initial Developer to use, reproduce, modify, display, perform, sublicense and distribute the Original Code (or portions thereof) with or without Modifications, and/or as part of a Larger Work; and

(b) under 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 Sections 2.2(a) and 2.2(b) are effective on the date Contributor first makes Commercial Use of the Covered Code.
- (d) Notwithstanding Section 2.2(b) above, no patent license is granted: 1) for any code that Contributor has deleted from the Contributor Version; 2) separate from the Contributor Version; 3) for infringements caused by: i) third party modifications of Contributor Version or ii) the combination of Modifications made by that Contributor with other software (except as part of the Contributor Version) or other devices; or 4) under Patent Claims infringed by Covered Code in the absence of Modifications made by that Contributor.

3. Distribution Obligations.

3.1. Application of License.

The Modifications which You create or to which You contribute are governed by the terms of this License, including without limitation Section 2.2. The Source Code version of Covered Code may be distributed only under the terms of this License or a future version of this License released under Section 6.1, and You must include a copy of this License with every copy of the Source Code You distribute. You may not offer or impose any terms on any Source Code version that alters or restricts the applicable version of this License or the recipients' rights hereunder. However, You may include an additional document offering the additional rights described in Section 3.5.

3.2. Availability of Source Code.

Any Modification which You create or to which You contribute must be made available in Source Code form under the terms of this License either on the same media as an Executable version or via an accepted Electronic Distribution Mechanism to anyone to whom you made an Executable version available; and if made available via Electronic Distribution Mechanism, must remain available for at least twelve (12) months after the date it initially became available, or at least six (6) months after a subsequent version of that particular Modification has been made available to such recipients. You are responsible for ensuring that the Source Code version remains available even if the Electronic Distribution Mechanism is maintained by a third party.

3.3. Description of Modifications.

You must cause all Covered Code to which You contribute to contain a file documenting the changes You made to create that Covered Code and the date of any change. You must include a prominent statement that the Modification is derived, directly or indirectly, from Original Code provided by the Initial Developer and including the name of the Initial Developer in (a) the Source Code, and (b) in any notice in an Executable version or related documentation in which You describe the origin or ownership of the Covered Code.

3.4. Intellectual Property Matters

- (a) Third Party Claims.

If Contributor has knowledge that a license under a third party's intellectual property rights is required to exercise the rights granted by such Contributor under Sections 2.1 or 2.2, Contributor must include a text file with the Source Code distribution titled "LEGAL" which describes the claim and the party making the claim in sufficient detail that a recipient will know whom to contact. If Contributor obtains such knowledge after the Modification is made available as described in Section 3.2, Contributor shall promptly modify the LEGAL file in all copies Contributor makes available thereafter and shall take other steps (such as notifying appropriate mailing lists or newsgroups) reasonably calculated to inform those who received the Covered Code that new knowledge has been obtained.

(b) Contributor APIs.

If Contributor's Modifications include an application programming interface and Contributor has knowledge of patent licenses which are reasonably necessary to implement that API, Contributor must also include this information in the LEGAL file.

(c) Representations.

Contributor represents that, except as disclosed pursuant to Section 3.4(a) above, Contributor believes that Contributor's Modifications are Contributor's original creation(s) and/or Contributor has sufficient rights to grant the rights conveyed by this License.

3.5. Required Notices.

You must duplicate the notice in Exhibit A in each file of the Source Code. If it is not possible to put such notice in a particular Source Code file due to its structure, then You must include such notice in a location (such as a relevant directory) where a user would be likely to look for such a notice. If You created one or more Modification(s) You may add your name as a Contributor to the notice described in Exhibit A. You must also duplicate this License in any documentation for the Source Code where You describe recipients' rights or ownership rights relating to Covered Code. You may choose to offer, and to charge a fee for, warranty, support, indemnity or liability obligations to one or more recipients of Covered Code. However, You may do so only on Your own behalf, and not on behalf of the Initial Developer or any Contributor. You must make it absolutely clear that any such warranty, support, indemnity or liability obligation is offered by You alone, and You hereby agree to indemnify the Initial Developer and every Contributor for any liability incurred by the Initial Developer or such Contributor as a result of warranty, support, indemnity or liability terms You offer.

3.6. Distribution of Executable Versions.

You may distribute Covered Code in Executable form only if the requirements of Section 3.1-3.5 have been met for that Covered Code, and if You include a notice stating that the Source Code version of the Covered Code is available under the terms of this License, including a description of how and where You have fulfilled the obligations of Section 3.2. The notice must be conspicuously included in any notice in an Executable version, related documentation or collateral in which You describe recipients' rights relating to the Covered Code. You may distribute the Executable version of Covered Code or ownership rights under a license of Your choice, which may contain terms different from this License, provided that You are in compliance with the terms of this License and that the license for the Executable version does not attempt to limit or alter the recipient's rights in the Source Code version from the rights set forth in this License. If You distribute the Executable version under a different license You must make it absolutely clear that any terms which differ from this License are offered by You alone, not by the Initial Developer or any Contributor. You hereby agree to indemnify the Initial Developer and every Contributor for any liability incurred by the Initial Developer or such Contributor as a result of any such terms You offer.

3.7. Larger Works.

You may create a Larger Work by combining Covered Code with other code not governed by the terms of this License and distribute the Larger Work as a single product. In such a case, You must make sure the requirements of this License are fulfilled for the Covered Code.

4. Inability to Comply Due to Statute or Regulation.

If it is impossible for You to comply with any of the terms of this License with respect to some or all of the Covered Code due to statute, judicial order, or regulation then You must: (a) comply with the terms of this License to the maximum extent possible; and (b) describe the limitations and the code they affect. Such description must be included in the LEGAL file described in Section 3.4 and must be included with all distributions of the Source Code. Except to the extent prohibited by statute or regulation, such description must be sufficiently detailed for a recipient of ordinary skill to be able to understand it.

5. Application of this License.

This License applies to code to which the Initial Developer has attached the notice in Exhibit A and to related Covered Code.

6. Versions of the License.

6.1. New Versions.

Netscape Communications Corporation ("Netscape") may publish revised and/or new versions of the License from time to time. Each version will be given a distinguishing version number.

6.2. Effect of New Versions.

Once Covered Code has been published under a particular version of the License, You may always continue to use it under the terms of that version. You may also choose to use such Covered Code under the terms of any subsequent version of the License published by Netscape. No one other than Netscape has the right to modify the terms applicable to Covered Code created under this License.

6.3. Derivative Works.

If You create or use a modified version of this License (which you may only do in order to apply it to code which is not already Covered Code governed by this License), You must (a) rename Your license so that the phrases "Mozilla", "MOZILLAPL", "MOZPL", "Netscape", "MPL", "NPL" or any confusingly similar phrase do not appear in your license (except to note that your licensed differs from this License) and (b) otherwise make it clear that Your version of the license contains terms which differ from the Mozilla Public License and Netscape Public License. (Filling in the name of the Initial Developer, Original Code or Contributor in the notice described in Exhibit A shall not of themselves be deemed to be modifications of this License.)

7. DISCLAIMER OF WARRANTY.

COVERED CODE IS PROVIDED UNDER THIS LICENSE ON AN "AS IS" BASIS, WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, WARRANTIES THAT THE COVERED CODE IS FREE OF DEFECTS, MERCHANTABLE, FIT FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE COVERED CODE IS WITH YOU. SHOULD ANY COVERED CODE PROVE DEFECTIVE IN ANY RESPECT, YOU (NOT THE INITIAL DEVELOPER OR ANY OTHER CONTRIBUTOR) ASSUME THE COST OF ANY NECESSARY SERVICING, REPAIR OR CORRECTION. THIS DISCLAIMER OF WARRANTY CONSTITUTES AN ESSENTIAL PART OF THIS LICENSE. NO USE OF ANY COVERED CODE IS AUTHORIZED HEREUNDER EXCEPT UNDER THIS DISCLAIMER.

8. TERMINATION.

8.1. This License and the rights granted hereunder will terminate automatically if You fail to comply with terms herein and fail to cure such breach within 30 days of becoming aware of the breach. All sublicenses to the Covered Code which are properly granted shall survive any termination of this License. Provisions which, by their nature, must remain in effect beyond the termination of this License shall survive.

8.2. If You initiate litigation by asserting a patent infringement claim (excluding declaratory judgment actions) against Initial Developer or a Contributor (the Initial Developer or Contributor against whom You file such action is referred to as "Participant") alleging that:

(a) such Participant's Contributor Version directly or indirectly infringes any patent, then any and all rights granted by such Participant to You under Sections 2.1 and/or 2.2 of this License shall, upon 60 days notice from Participant terminate prospectively, unless if within 60 days after receipt of notice You either: (i) agree in writing to pay Participant a mutually agreeable reasonable royalty for Your past and future use of Modifications made by such Participant, or (ii) withdraw Your litigation claim with respect to the Contributor Version against such Participant. If within 60 days of notice, a reasonable royalty and payment arrangement are not mutually agreed upon in writing by the parties or the litigation claim is not withdrawn, the rights granted by Participant to You under Sections 2.1 and/or 2.2 automatically terminate at the expiration of the 60 day notice period specified above.

(b) any software, hardware, or device, other than such Participant's Contributor Version, directly or indirectly infringes any patent, then any rights granted to You by such Participant under Sections 2.1(b) and 2.2(b) are revoked effective as of the date You first made, used, sold, distributed, or had made, Modifications made by that Participant.

8.3. If You assert a patent infringement claim against Participant alleging that such Participant's Contributor Version directly or indirectly infringes any patent where such claim is resolved (such as by license or settlement) prior to the initiation of patent infringement litigation, then the reasonable value of the licenses granted by such Participant under Sections 2.1 or 2.2 shall be taken into account in determining the amount or value of any payment or license.

8.4. In the event of termination under Sections 8.1 or 8.2 above, all end user license agreements (excluding distributors and resellers) which have been validly granted by You or any distributor hereunder prior to termination shall survive termination.

9. LIMITATION OF LIABILITY.

UNDER NO CIRCUMSTANCES AND UNDER NO LEGAL THEORY, WHETHER TORT (INCLUDING NEGLIGENCE), CONTRACT, OR OTHERWISE, SHALL YOU, THE INITIAL DEVELOPER, ANY OTHER CONTRIBUTOR, OR ANY DISTRIBUTOR OF COVERED CODE, OR ANY SUPPLIER OF ANY OF SUCH PARTIES, BE LIABLE TO ANY PERSON FOR ANY INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES OF ANY CHARACTER INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF GOODWILL, WORK STOPPAGE, COMPUTER FAILURE OR MALFUNCTION, OR ANY AND ALL OTHER COMMERCIAL DAMAGES OR LOSSES, EVEN IF SUCH PARTY SHALL HAVE BEEN INFORMED OF THE POSSIBILITY OF SUCH DAMAGES. THIS LIMITATION OF LIABILITY SHALL NOT APPLY TO LIABILITY FOR DEATH OR PERSONAL INJURY RESULTING FROM SUCH PARTY'S NEGLIGENCE TO THE EXTENT APPLICABLE LAW PROHIBITS SUCH LIMITATION. SOME JURISDICTIONS DO NOT ALLOW THE EXCLUSION OR LIMITATION OF INCIDENTAL OR CONSEQUENTIAL DAMAGES, SO THIS EXCLUSION AND LIMITATION MAY NOT APPLY TO YOU.

10. U.S. GOVERNMENT END USERS.

The Covered Code is a "commercial item," as that term is defined in 48 C.F.R. 2.101 (Oct. 1995), consisting of "commercial computer software" and "commercial computer software documentation," as such terms are used in 48 C.F.R. 12.212 (Sept. 1995). Consistent with 48 C.F.R. 12.212 and 48 C.F.R. 227.7202-1 through 227.7202-4 (June 1995), all U.S. Government End Users acquire Covered Code with only those rights set forth herein.

11. MISCELLANEOUS.

This License represents the complete agreement concerning subject matter hereof. If any provision of this License is held to be unenforceable, such provision shall be reformed only to the extent necessary to make it enforceable. This License shall be governed by California law provisions (except to the extent applicable law, if any, provides otherwise), excluding its conflict-of-law provisions. With respect to disputes in which at least one party is a citizen of, or an entity chartered or registered to do business in the United States of America, any litigation relating to this License shall be subject to the jurisdiction of the Federal Courts of the Northern District of California, with venue lying in Santa Clara County, California, with the losing party responsible for costs, including without limitation, court costs and reasonable attorneys' fees and expenses. The application of the United Nations Convention on Contracts for the International Sale of Goods is expressly excluded. Any law or regulation which provides that the language of a contract shall be construed against the drafter shall not apply to this License.

12. RESPONSIBILITY FOR CLAIMS.

As between Initial Developer and the Contributors, each party is responsible for claims and damages arising, directly or indirectly, out of its utilization of rights under this License and You agree to work with Initial Developer and Contributors to distribute such responsibility on an equitable basis. Nothing herein is intended or shall be deemed to constitute any admission of liability.

13. MULTIPLE-LICENSED CODE.

Initial Developer may designate portions of the Covered Code as "Multiple-Licensed". "Multiple-Licensed" means that the Initial Developer permits you to utilize portions of the Covered Code under Your choice of the NPL or the alternative licenses, if any, specified by the Initial Developer in the file described in Exhibit A.

EXHIBIT A -Mozilla Public License.

`` The contents of this file are subject to the Mozilla Public License Version 1.1 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at <http://www.mozilla.org/MPL/>

Software distributed under the License is distributed on an "AS IS" basis, WITHOUT WARRANTY OF ANY KIND, either express or implied. See the License for the specific language governing rights and limitations under the License.

The Original Code is _____.

The Initial Developer of the Original Code is _____.

Portions created by _____ are Copyright (C) _____
_____. All Rights Reserved.

Contributor(s): _____.

Alternatively, the contents of this file may be used under the terms of the _____ license (the "[_____] License"), in which case the provisions of [_____] License are applicable instead of those above. If you wish to allow use of your version of this file only under the terms of the [_____] License and not to allow others to use your version of this file under the MPL, indicate your decision by deleting the provisions above and replace them with the notice and other provisions required by the [_____] License. If you do not delete the provisions above, a recipient may use your version of this file under either the MPL or the [_____] License."

[NOTE: The text of this Exhibit A may differ slightly from the text of the notices in the Source Code files of the Original Code. You should use the text of this Exhibit A rather than the text found in the Original Code Source Code for Your Modifications.]

CDDL 1.1

This license applies to the following open source components:

- jaxb-api-2.3.1.jar - <https://mvnrepository.com/artifact/javax.xml.bind/jaxb-api/2.3.1>
Source code is available at: <https://github.com/slcorp/jaxb-api-2.3.1-sources>
- jaxb-runtime-2.3.1.jar - <https://mvnrepository.com/artifact/org.glassfish.jaxb/jaxb-runtime/2.3.1>
Source code is available at: <https://github.com/slcorp/jaxb-runtime-2.3.1-sources>
- javax.activation-api-1.2.0.jar – <https://mvnrepository.com/artifact/javax.activation/javax.activation-api/1.2.0>
Source code is available at: <https://github.com/slcorp/javax.activation-api-1.2.0-sources>
- jaxb-core-2.3.0.1.jar - <https://mvnrepository.com/artifact/com.sun.xml.bind/jaxb-core/2.3.0.1>
Source code is available at: <https://github.com/slcorp/jaxb-core-2.3.0.1-sources>

COMMON DEVELOPMENT AND DISTRIBUTION LICENSE (CDDL) Version 1.1

1. Definitions.

1.1. "Contributor" means each individual or entity that creates or contributes to the creation of Modifications.

1.2. "Contributor Version" means the combination of the Original Software, prior Modifications used by a Contributor (if any), and the Modifications made by that particular Contributor.

1.3. "Covered Software" means (a) the Original Software, or (b) Modifications, or (c) the combination of files containing Original Software with files containing Modifications, in each case including portions thereof.

1.4. "Executable" means the Covered Software in any form other than Source Code.

1.5. "Initial Developer" means the individual or entity that first makes Original Software available under this License.

1.6. "Larger Work" means a work which combines Covered Software or portions thereof with code not governed by the terms of this License.

1.7. "License" means this document.

1.8. "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 the Source Code and Executable form of any of the following:

- A. Any file that results from an addition to, deletion from or modification of the contents of a file containing Original Software or previous Modifications;
- B. Any new file that contains any part of the Original Software or previous Modification; or
- C. Any new file that is contributed or otherwise made available under the terms of this License.

1.10. "Original Software" means the Source Code and Executable form of computer software code that is originally released under this License.

1.11. "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.

1.12. "Source Code" means (a) the common form of computer software code in which modifications are made and (b) associated documentation included in or with such code.

1.13. "You" (or "Your") means an individual or a legal entity exercising rights under, and complying with all of the terms of, this License. For legal entities, "You" includes any entity which controls, is controlled by, or is under common control with You. For purposes of this definition, "control" means (a) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (b) ownership of more than fifty percent (50%) of the outstanding shares or beneficial ownership of such entity.

2. License Grants.

2.1. The Initial Developer Grant.

Conditioned upon Your compliance with Section 3.1 below and subject to third party intellectual property claims, the Initial Developer hereby grants You a world-wide, royalty-free, non-exclusive license:

(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 Software (or portions thereof), with or without Modifications, and/or as part of a Larger Work; and

(b) under Patent Claims infringed by the making, using or selling of Original Software, to make, have made, use, practice, sell, and offer for sale, and/or otherwise dispose of the Original Software (or portions thereof).

(c) The licenses granted in Sections 2.1(a) and (b) are effective on the date Initial Developer first distributes or otherwise makes the Original Software available to a third party 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 Software, or (2) for infringements caused by: (i) the modification of the Original Software, or (ii) the combination of the Original Software with other software or devices.

2.2. Contributor Grant.

Conditioned upon Your compliance with Section 3.1 below and 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 Software and/or as part of a Larger Work; and (b) under Patent Claims infringed by the making, using, or selling of Modifications made by that Contributor either alone and/or in combination with its Contributor Version (or portions of such combination), to make, use, sell, offer for sale, have made, and/or otherwise dispose of: (1) Modifications made by that Contributor (or portions thereof); and (2) the combination of Modifications made by that Contributor with its Contributor Version (or portions of such combination).

(c) The licenses granted in Sections 2.2(a) and 2.2(b) are effective on the date Contributor first distributes or otherwise makes the Modifications available to a third party.

(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) for infringements caused by: (i) third party modifications of Contributor Version, or (ii) the combination of Modifications made by that Contributor with other software (except as part of the Contributor Version) or other devices; or (3) under Patent Claims infringed by Covered Software in the absence of Modifications made by that Contributor.

3. Distribution Obligations.

3.1. Availability of Source Code.

Any Covered Software that You distribute or otherwise make available in Executable form must also be made available in Source Code form and that Source Code form must be distributed only under the terms of this License. You must include a copy of this License with every copy of the Source Code form of the Covered Software You distribute or otherwise make available. You must inform recipients of any such Covered Software in Executable form as to how they can obtain such Covered Software in Source Code form in a reasonable manner on or through a medium customarily used for software exchange.

3.2. Modifications.

The Modifications that You create or to which You contribute are governed by the terms of this License. You represent that You believe Your Modifications are Your original creation(s) and/or You have sufficient rights to grant the rights conveyed by this License.

3.3. Required Notices.

You must include a notice in each of Your Modifications that identifies You as the Contributor of the Modification. You may not remove or alter any copyright, patent or trademark notices contained within the Covered Software, or any notices of licensing or any descriptive text giving attribution to any Contributor or the Initial Developer.

3.4. Application of Additional Terms.

You may not offer or impose any terms on any Covered Software in Source Code form that alters or restricts the applicable version of this License or the recipients' rights hereunder. You may choose to offer, and to charge a fee for, warranty, support, indemnity or liability obligations to one or more recipients of Covered Software. However, you may do so only on Your own behalf, and not on behalf of the Initial Developer or any Contributor. You must make it absolutely clear that any such warranty, support, indemnity or liability obligation is offered by You alone, and You hereby agree to indemnify the Initial Developer and every Contributor for any liability incurred by the Initial Developer or such Contributor as a result of warranty, support, indemnity or liability terms You offer.

3.5. Distribution of Executable Versions.

You may distribute the Executable form of the Covered Software under the terms of this License or under the terms of a license of Your choice, which may contain terms different from this License, provided that You are in compliance with the terms of this License and that the license for the Executable form does not attempt to limit or alter the recipient's rights in the Source Code form from the rights set forth in this License. If You distribute the Covered Software in Executable form 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 Contributor. You hereby agree to indemnify the Initial Developer and every Contributor for any liability incurred by the Initial Developer or such Contributor as a result of any such terms You offer.

3.6. Larger Works.

You may create a Larger Work by combining Covered Software 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 Software.

4. Versions of the License.

4.1. New Versions.

Oracle is the initial license steward and may publish revised and/or new versions of this License from time to time. Each version will be given a distinguishing version number. Except as provided in Section 4.3, no one other than the license steward has the right to modify this License.

4.2. Effect of New Versions.

You may always continue to use, distribute or otherwise make the Covered Software available under the terms of the version of the License under which You originally received the Covered Software. If the Initial Developer includes a notice in the Original Software prohibiting it from being distributed or otherwise made available under any subsequent version of the License, You must distribute and make the Covered Software available under the terms of the version of the License under which You originally received the Covered Software. Otherwise, You may also choose to use, distribute or otherwise make the Covered Software available under the terms of any subsequent version of the License published by the license steward.

4.3. Modified Versions.

When You are an Initial Developer and You want to create a new license for Your Original Software, You may create and use a modified version of this License if You: (a) rename the license and remove any references to the name of the license steward (except to note that the license differs from this License); and (b) otherwise make it clear that the license contains terms which differ from this License.

5. DISCLAIMER OF WARRANTY.

COVERED SOFTWARE 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 SOFTWARE IS FREE OF DEFECTS, MERCHANTABLE, FIT FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE COVERED SOFTWARE IS WITH YOU. SHOULD ANY COVERED SOFTWARE 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 SOFTWARE IS AUTHORIZED HEREUNDER EXCEPT UNDER THIS DISCLAIMER.

6. TERMINATION.

6.1. This License and the rights granted hereunder will terminate automatically if You fail to comply with terms herein and fail to cure such breach within 30 days of becoming aware of the breach. Provisions which, by their nature, must remain in effect beyond the termination of this License shall survive.

6.2. If You assert a patent infringement claim (excluding declaratory judgment actions) against Initial Developer or a Contributor (the Initial Developer or Contributor against whom You assert such claim is referred to as "Participant") alleging that the Participant Software (meaning the Contributor Version where the Participant is a Contributor or the Original Software where the Participant is the Initial Developer) directly or indirectly infringes any patent, then any and all rights granted directly or indirectly to You by such Participant, the Initial Developer (if the Initial Developer is not the Participant) and all Contributors under Sections 2.1 and/or 2.2 of this License shall, upon 60 days notice from Participant terminate prospectively and automatically at the expiration of such 60 day notice period, unless if within such 60 day period You withdraw Your claim with respect to the Participant Software against such Participant either unilaterally or pursuant to a written agreement with Participant.

6.3. If You assert a patent infringement claim against Participant alleging that the Participant Software directly or indirectly infringes any patent where such claim is resolved (such as by license or settlement) prior to the initiation of patent infringement litigation, then the reasonable value of the licenses granted by such Participant under Sections 2.1 or 2.2 shall be taken into account in determining the amount or value of any payment or license.

6.4. In the event of termination under Sections 6.1 or 6.2 above, all end user licenses that have been validly granted by You or any distributor hereunder prior to termination (excluding licenses granted to You by any distributor) shall survive termination.

7. 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 SOFTWARE, 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.

8. U.S. GOVERNMENT END USERS.

The Covered Software is a "commercial item," as that term is defined in 48 C.F.R. 2.101 (Oct. 1995), consisting of "commercial computer software" (as that term is defined at 48 C.F.R. § 252.227-7014(a)(1)) and "commercial computer software documentation" as such terms are used in 48 C.F.R. 12.212 (Sept. 1995). Consistent with 48 C.F.R. 12.212 and 48 C.F.R. 227.7202-1 through 227.7202-4 (June 1995), all U.S. Government End Users acquire Covered Software with only those rights set forth herein. This U.S. Government Rights clause is in lieu of, and supersedes, any other FAR, DFAR, or other clause or provision that addresses Government rights in computer software under this License.

9. MISCELLANEOUS.

This License represents the complete agreement concerning subject matter hereof. If any provision of this License is held to be unenforceable, such provision shall be reformed only to the extent necessary to make it enforceable. This License shall be governed by the law of the jurisdiction specified in a notice contained within the Original Software (except to the extent applicable law, if any, provides otherwise), excluding such jurisdiction's conflict-of-law provisions. Any litigation relating to this License shall be subject to the jurisdiction of the courts located in the jurisdiction and venue specified in a notice contained within the Original Software, with the losing party responsible for costs, including, without limitation, court costs and reasonable attorneys' fees and expenses. The application of the United Nations Convention on Contracts for the International Sale of Goods is expressly excluded. Any law or regulation which provides that the language of a contract shall be construed against the drafter shall not apply to this License. You agree that You alone are responsible for compliance with the United States export administration regulations (and the export control laws and regulation of any other countries) when You use, distribute or otherwise make available any Covered Software.

10. 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.

NOTICE PURSUANT TO SECTION 9 OF THE COMMON DEVELOPMENT AND DISTRIBUTION LICENSE (CDDL)

The code released under the CDDL shall be governed by the laws of the State of California (excluding conflict-of-law provisions). Any litigation relating to this License shall be subject to the jurisdiction of the Federal Courts of the Northern District of California and the state courts of the State of California, with venue lying in Santa Clara County, California.

=====

****MD Datejs**

Copyright © 2006-2010 Coolite Inc.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

=====

****jQuery**

Copyright © 2009 John Resig

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

=====

** JCalendar 1.3.2

This product uses JCalendar 1.3.2. JCalendar is distributed pursuant to the terms of the Lesser General Public License. The source code for the JCalendar may be obtained from <http://www.toedter.com/en/jcalendar/index.html>

=====

** BrowserLauncher2 1.3

This product uses BrowserLauncher 1.3 and is distributed pursuant to the terms of the Lesser General Public License. The source code for BrowserLauncher2 1.3 can be obtained from: <http://browserlaunch2.sourceforge.net/>

GNU LESSER GENERAL PUBLIC LICENSE

Version 2.1, February 1999

Copyright (C) 1991, 1999 Free Software Foundation, Inc.

59 Temple Place, Suite 330, Boston, MA 02111-1307 USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

[This is the first released version of the Lesser GPL. It also counts as the successor of the GNU Library Public License, version 2, hence the version number 2.1.]

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public Licenses are intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users.

This license, the Lesser General Public License, applies to some specially designated software packages--typically libraries--of the Free Software Foundation and other authors who decide to use it. You can use it too, but we suggest you first think carefully about whether this license or the ordinary General Public License is the better strategy to use in any particular case, based on the explanations below.

When we speak of free software, we are referring to freedom of use, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish); that you receive source code or can get it if you want it; that you can change the software and use pieces of it in new free programs; and that you are informed that you can do these things.

To protect your rights, we need to make restrictions that forbid distributors to deny you these rights or to ask you to surrender these rights. These restrictions translate to certain responsibilities for you if you distribute copies of the library or if you modify it.

For example, if you distribute copies of the library, whether gratis or for a fee, you must give the recipients all the rights that we gave you. You must make sure that they, too, receive or can get the source code. If you link other code with the library, you must provide complete object files to the recipients, so that they can relink them with the library after making changes to the library and recompiling it. And you must show them these terms so they know their rights.

We protect your rights with a two-step method: (1) we copyright the library, and (2) we offer you this license, which gives you legal permission to copy, distribute and/or modify the library.

To protect each distributor, we want to make it very clear that there is no warranty for the free library. Also, if the library is modified by someone else and passed on, the recipients should know that what they have is not the original version, so that the original author's reputation will not be affected by problems that might be introduced by others.

Finally, software patents pose a constant threat to the existence of any free program. We wish to make sure that a company cannot effectively restrict the users of a free program by obtaining a restrictive license from a patent holder. Therefore, we insist that any patent license obtained for a version of the library must be consistent with the full freedom of use specified in this license.

Most GNU software, including some libraries, is covered by the ordinary GNU General Public License. This license, the GNU Lesser General Public License, applies to certain designated libraries, and is quite different from the ordinary General Public License. We use this license for certain libraries in order to permit linking those libraries into non-free programs.

When a program is linked with a library, whether statically or using a shared library, the combination of the two is legally speaking a combined work, a derivative of the original library. The ordinary General Public License therefore permits such linking only if the entire combination fits its criteria of freedom. The Lesser General Public License permits more lax criteria for linking other code with the library.

We call this license the "Lesser" General Public License because it does Less to protect the user's freedom than the ordinary General Public License. It also provides other free software developers Less of an advantage over competing non-free programs. These disadvantages are the reason we use the ordinary General Public License for many libraries. However, the Lesser license provides advantages in certain special circumstances.

For example, on rare occasions, there may be a special need to encourage the widest possible use of a certain library, so that it becomes a de-facto standard. To achieve this, non-free programs must be allowed to use the library. A more frequent case is that a free library does the same job as widely used non-free libraries. In this case, there is little to gain by limiting the free library to free software only, so we use the Lesser General Public License.

In other cases, permission to use a particular library in non-free programs enables a greater number of people to use a large body of free software. For example, permission to use the GNU C Library in non-free programs enables many more people to use the whole GNU operating system, as well as its variant, the GNU/Linux operating system.

Although the Lesser General Public License is Less protective of the users' freedom, it does ensure that the user of a program that is linked with the Library has the freedom and the wherewithal to run that program using a modified version of the Library.

The precise terms and conditions for copying, distribution and modification follow. Pay close attention to the difference between a "work based on the library" and a "work that uses the library". The former contains code derived from the library, whereas the latter must be combined with the library in order to run.

TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License Agreement applies to any software library or other program which contains a notice placed by the copyright holder or other authorized party saying it may be distributed under the terms of this Lesser General Public License (also called "this License"). Each licensee is addressed as "you".

A "library" means a collection of software functions and/or data prepared so as to be conveniently linked with application programs (which use some of those functions and data) to form executables.

The "Library", below, refers to any such software library or work which has been distributed under these terms. A "work based on the Library" means either the Library or any derivative work under copyright law: that is to say, a work containing the Library or a portion of it, either verbatim or with modifications and/or translated straightforwardly into another language. (Hereinafter, translation is included without limitation in the term "modification".)

"Source code" for a work means the preferred form of the work for making modifications to it. For a library, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the library.

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running a program using the Library is not restricted, and output from such a program is covered only if its contents constitute a work based on the Library (independent of the use of the Library in a tool for writing it). Whether that is true depends on what the Library does and what the program that uses the Library does.

1. You may copy and distribute verbatim copies of the Library's complete source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and distribute a copy of this License along with the Library.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Library or any portion of it, thus forming a work based on the Library, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

- a) The modified work must itself be a software library.
- b) You must cause the files modified to carry prominent notices stating that you changed the files and the date of any change.
- c) You must cause the whole of the work to be licensed at no charge to all third parties under the terms of this License.

d) If a facility in the modified Library refers to a function or a table of data to be supplied by an application program that uses the facility, other than as an argument passed when the facility is invoked, then you must make a good faith effort to ensure that, in the event an application does not supply such function or table, the facility still operates, and performs whatever part of its purpose remains meaningful.

(For example, a function in a library to compute square roots has a purpose that is entirely well-defined independent of the application. Therefore, Subsection 2d requires that any application-supplied function or table used by this function must be optional: if the application does not supply it, the square root function must still compute square roots.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Library, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Library, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Library.

In addition, mere aggregation of another work not based on the Library with the Library (or with a work based on the Library) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may opt to apply the terms of the ordinary GNU General Public License instead of this License to a given copy of the Library. To do this, you must alter all the notices that refer to this License, so that they refer to the ordinary GNU General Public License, version 2, instead of to this License. (If a newer version than version 2 of the ordinary GNU General Public License has appeared, then you can specify that version instead if you wish.) Do not make any other change in these notices.

Once this change is made in a given copy, it is irreversible for that copy, so the ordinary GNU General Public License applies to all subsequent copies and derivative works made from that copy.

This option is useful when you wish to copy part of the code of the Library into a program that is not a library.

4. You may copy and distribute the Library (or a portion or derivative of it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange.

If distribution of object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place satisfies the requirement to distribute the source code, even though third parties are not compelled to copy the source along with the object code.

5. A program that contains no derivative of any portion of the Library, but is designed to work with the Library by being compiled or linked with it, is called a "work that uses the Library". Such a work, in isolation, is not a derivative work of the Library, and therefore falls outside the scope of this License.

However, linking a "work that uses the Library" with the Library creates an executable that is a derivative of the Library (because it contains portions of the Library), rather than a "work that uses the library". The executable is therefore covered by this License. Section 6 states terms for distribution of such executables.

When a "work that uses the Library" uses material from a header file that is part of the Library, the object code for the work may be a derivative work of the Library even though the source code is not. Whether this is true is especially significant if the work can be linked without the Library, or if the work is itself a library. The threshold for this to be true is not precisely defined by law.

If such an object file uses only numerical parameters, data structure layouts and accessors, and small macros and small inline functions (ten lines or less in length), then the use of the object file is unrestricted, regardless of whether it is legally a derivative work. (Executables containing this object code plus portions of the Library will still fall under Section 6.)

Otherwise, if the work is a derivative of the Library, you may distribute the object code for the work under the terms of Section 6. Any executables containing that work also fall under Section 6, whether or not they are linked directly with the Library itself.

6. As an exception to the Sections above, you may also combine or link a "work that uses the Library" with the Library to produce a work containing portions of the Library, and distribute that work under terms of your choice, provided that the terms permit modification of the work for the customer's own use and reverse engineering for debugging such modifications.

You must give prominent notice with each copy of the work that the Library is used in it and that the Library and its use are covered by this License. You must supply a copy of this License. If the work during execution displays copyright notices, you must include the copyright notice for the Library among them, as well as a reference directing the user to the copy of this License. Also, you must do one of these things:

- a) Accompany the work with the complete corresponding machine-readable source code for the Library including whatever changes were used in the work (which must be distributed under Sections 1 and 2 above); and, if the work is an executable linked with the Library, with the complete machine-readable "work that uses the Library", as object code and/or source code, so that the user can modify the Library and then relink to produce a modified executable containing the modified Library. (It is understood that the user who changes the contents of definitions files in the Library will not necessarily be able to recompile the application to use the modified definitions.)
- b) Use a suitable shared library mechanism for linking with the Library. A suitable mechanism is one that (1) uses at run time a copy of the library already present on the user's computer system, rather than copying library functions into the executable, and (2) will operate properly with a modified version of the library, if the user installs one, as long as the modified version is interface-compatible with the version that the work was made with.
- c) Accompany the work with a written offer, valid for at least three years, to give the same user the materials specified in Subsection 6a, above, for a charge no more than the cost of performing this distribution.
- d) If distribution of the work is made by offering access to copy from a designated place, offer equivalent access to copy the above specified materials from the same place.
- e) Verify that the user has already received a copy of these materials or that you have already sent this user a copy.

For an executable, the required form of the "work that uses the Library" must include any data and utility programs needed for reproducing the executable from it. However, as a special exception, the materials to be distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

It may happen that this requirement contradicts the license restrictions of other proprietary libraries that do not normally accompany the operating system. Such a contradiction means you cannot use both them and the Library together in an executable that you distribute.

7. You may place library facilities that are a work based on the Library side-by-side in a single library together with other library facilities not covered by this License, and distribute such a combined library, provided that the separate distribution of the work based on the Library and of the other library facilities is otherwise permitted, and provided that you do these two things:

- a) Accompany the combined library with a copy of the same work based on the Library, uncombined with any other library facilities. This must be distributed under the terms of the Sections above.
- b) Give prominent notice with the combined library of the fact that part of it is a work based on the Library, and explaining where to find the accompanying uncombined form of the same work.

8. You may not copy, modify, sublicense, link with, or distribute the Library except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, link with, or distribute the Library is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

9. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Library or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Library (or any work based on the Library), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Library or works based on it.

10. Each time you redistribute the Library (or any work based on the Library), the recipient automatically receives a license from the original licensor to copy, distribute, link with or modify the Library subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties with this License.

11. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Library at all. For example, if a patent license would not permit royalty-free redistribution of the Library by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Library.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply, and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

12. If the distribution and/or use of the Library is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Library under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

13. The Free Software Foundation may publish revised and/or new versions of the Lesser General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Library specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Library does not specify a license version number, you may choose any version ever published by the Free Software Foundation.

14. If you wish to incorporate parts of the Library into other free programs whose distribution conditions are incompatible with these, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

15. BECAUSE THE LIBRARY IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE LIBRARY, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE LIBRARY "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE LIBRARY IS WITH YOU. SHOULD THE LIBRARY PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

16. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE LIBRARY AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE LIBRARY (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE LIBRARY TO OPERATE WITH ANY OTHER SOFTWARE), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

How to Apply These Terms to Your New Libraries

If you develop a new library, and you want it to be of the greatest possible use to the public, we recommend making it free software that everyone can redistribute and change. You can do so by permitting redistribution under these terms (or, alternatively, under the terms of the ordinary General Public License).

To apply these terms, attach the following notices to the library. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

one line to give the library's name and an idea of what it does.

Copyright (C) year name of author

This library is free software; you can redistribute it and/or modify it under the terms of the GNU Lesser General Public

License as published by the Free Software Foundation; either version 2.1 of the License, or (at your option) any later version.

This library is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU Lesser General Public License for more details.

You should have received a copy of the GNU Lesser General Public License along with this library; if not, write to the Free Software Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA

Also add information on how to contact you by electronic and paper mail.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a "copyright disclaimer" for the library, if necessary. Here is a sample; alter the names:

Yoyodyne, Inc., hereby disclaims all copyright interest in the library `Frob' (a library for tweaking knobs) written by James Random Hacker.

signature of Ty Coon, 1 April 1990

Ty Coon, President of Vice

That's all there is to it!

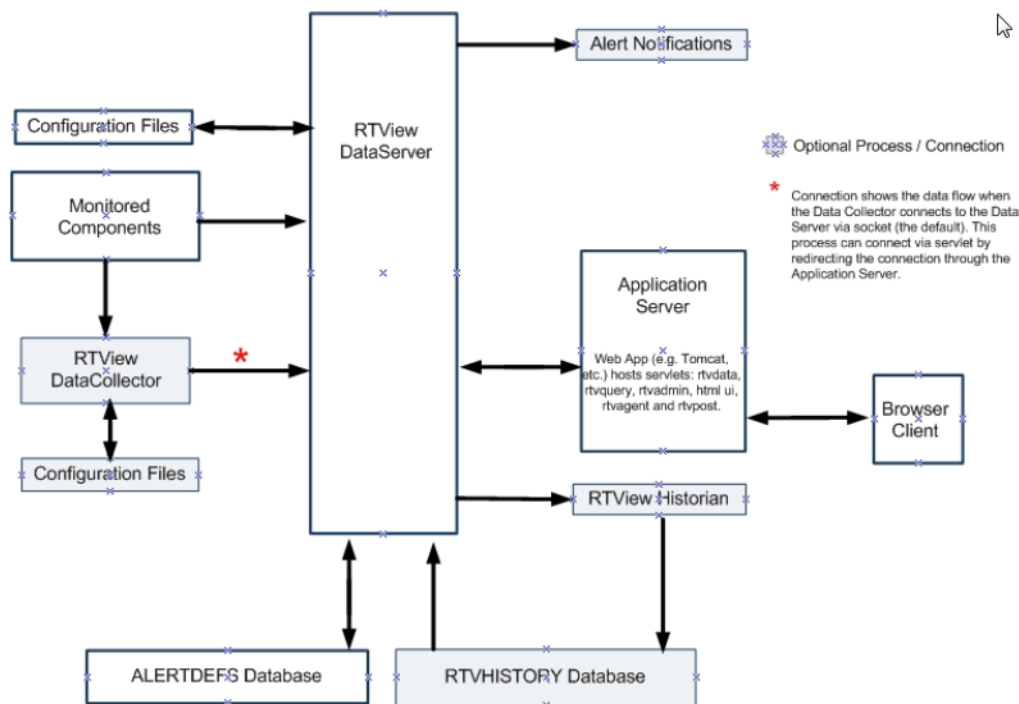
APPENDIX E Security Configuration

This section provides details for securing a direct connection Solace PubSub+ Monitor deployment. This section contains:

- ["Introduction"](#)
- ["Data Server"](#)
- ["HTML UI"](#)
- ["Data Collector"](#)
- ["Configuration Application"](#)
- ["Configuration Files"](#)
- ["Historian"](#)
- ["Database"](#)
- ["Application Servers"](#)
- ["Monitored Components"](#)
- ["Security Summary"](#)

Introduction

The following Solace PubSub+ Monitoring diagram shows how data flows through the Solace PubSub+ Monitor deployment. The Data Server connects to the Monitored Components to collect metric data which it stores in local caches and uses to generate alerts based on the enabled and threshold settings in the ALERTDEFS database and optionally execute user defined alert notifications.



In cases where the data collection needs to be distributed, one or more Data Collectors can be deployed to connect to the **"Monitored Components"** and forward the collected data to the **"Data Server"**.

The **"HTML UI"** is a browser-based user interface that shows metric and alert data from the Data Server and also allow the user to enable, disable and set thresholds on alerts.

The **"Historian"** is an optional process that stores historical metric and alert data to the RTVHISTORY database. When the Historian is enabled, the Data Server will query historical data from the RTVHISTORY on startup to populate in-memory history and also any time the HTML UI request history data that is older than the data in the in-memory history.

The **"Configuration Application"** is a browser based application for configuring the RTView processes. It connects to the Data Server to read and write **"Configuration Files"**.

The next sections provide a more detailed description of each process.

Data Server

The Data Server gathers and caches the data from the applications being monitored and also hosts the alerts for that data. Because the Data Server can exist behind firewalls, it simplifies and strengthens the secured delivery of information to clients beyond the firewall. The Data Server serves metric and alert data to the Historian via socket on port **4178** and receives data via socket from the optional Data Collector on port **4172**. It also serves metrics and alert data to the HTML UI via the `rtvquery` servlet which connects via socket on port **4178**.

The [“Historian”](#) runs in the same directory as the Data Server, while the optional Data Collector(s) typically run in a different directory or a different system. By default, socket connections to the Data Server are unsecured. The Data Server supports secure socket connections (SSL) without certificates. Once the Data Server has been configured for secure sockets, the History and all `rtvdata` and `rtvquery` servlets will connect via secure socket.

The [“HTML UI”](#) connects to the Data Server via the `rtvquery` servlet. See the [“HTML UI”](#) section in this document for information on how to enable authentication in the HTML UI and `rtvquery` servlets. The `rtvquery` servlet will connect to the Data Server via secure socket if the Data Server is configured for SSL sockets.

The [“Data Collector”](#) is an optional process that is used to distribute connections to Monitored Components Data Collectors instead of having the Data Server connect to each component to be monitored directly. This process collects data from Monitored Components and forwards it to the Data Server via socket or the `rtvagent` servlet. See the [“Data Collector”](#) section for information on securing this connection.

The [“Configuration Application”](#) connects to the Data Server via the `rtvadmin` servlet to read and write properties files. The `rtvadmin` servlet connects to the Data Server via socket on port **4178**. See the [“Configuration Application”](#) section below for information about servlet authentication. The `rtvadmin` servlet will connect via secure socket if the Data Server is configured for SSL sockets.

If the Historian is enabled, the Data Server connects to the RTVHISTORY database on startup to read initial cache history data and if the HTML UI request history data older than the in memory cache history. It also connects to the ALERTDEFS database to query and set alert thresholds. See the [“Database”](#) section below for more information.

The Data Server optionally executes alert notifications based on user settings. Since the notification actions are user defined, security must be determined by the user.

The Data Server opens a JMX port on **4168** to enable monitoring. By default, the JMX port is not secured. See [“Monitored Components”](#) for information on securing this connection.

HTML UI

This interface is implemented in HTML and is deployed as a servlet, **rtview-solmon**, which is configured by default to use HTTP authentication. Browser clients connect via HTTP or HTTPS depending on the Application Server configuration. For secure deployments, it should be configured to use HTTPS since HTTP authentication does not encrypt user credentials. The HTML UI sends data requests to the `rtvquery` servlet which connects to the Data Server via socket. By default, this socket is unsecured, but the `rtvquery` servlet will connect to the Data Server via secure socket if the Data Server is configured for SSL sockets.

By default, the `rtvquery` servlet *is* configured for authentication, but you can unsecure it by doing the following:

- `cd` to `rtvapm` and run `rtvapm_init`
- `cd` to `projects\rtview-server`
- edit `update_wars.bat` (or `.sh`) to remove the line that sets `SECURE=-secure`
- run `update_wars.bat` (or `.sh`)

Data Collector

This process is optional and is used to distribute connections to Monitored Components Data Collectors instead of having the Data Server connect to each component to be monitored directly. This process collects data from Monitored Components and forwards it to the Data Server via socket on port **4172** or the `rtvagent` servlet. In the RTView Configuration Application Data Server **COLLECTOR** tab, the **Target** definition determines whether data is sent to a socket or a servlet. If the URL for the target is **host:port**, it will be sent via socket which is not secured by default. This socket can be secured via SSL by specifying the following property on the **CUSTOM PROPERTIES** tab in the Configuration Application of each receiver Data Server:

Property Name: `sl.rtvview.rtvagent.ssl`

Property Value: `true`

Property Filter: `collector`

If the URL is the receiver's `rtvagent` servlet it will send data to that `rtvagent` servlet which will connect via socket to the Data Server on port **4172** which can be secured via SSL as described above. While the `rtvagent` servlet cannot be configured for authentication, Tomcat access filters can be used to restrict access and it can be deployed on HTTPS. While the Data Collector typically does not have data clients, it accepts data requests via socket on port **4176** which can be secured as described in the [“Data Server”](#) section. The Data Collector also opens JMX on port **4166** for monitoring. By default, the JMX port is not secured. See [“Monitored Components”](#) for information on securing this connection.

Configuration Application

The Configuration Application connects to the Data Server via the `rtvadmin` servlet which is configured with HTTP authentication. It should be run on HTTPS since user credentials are not encrypted. Passwords saved by the configuration application are scrambled except in the case where they are added in the **CUSTOM PROPERTIES** section. The `rtvadmin` servlet connects to the Data Server via socket. By default, this socket is unsecured, but the `rtvadmin` servlet will connect to the [“Data Server”](#) via secure socket if the Data Server is configured for SSL sockets.

Configuration Files

Configuration (**.properties**) files are stored on the file system and read by all RTView processes to control configuration. Additionally, the Configuration Application writes these files, scrambling all connection and database passwords. Passwords entered in the **CUSTOM PROPERTIES** tab are not scrambled.

Historian

The Historian connects to the [“Data Server”](#) via socket and saves cache history to a database via JDBC. This process is optional and the user can configure which data will be saved. By default, the socket connection is unsecured, but the Historian will connect via secure socket if the Data Server is configured for SSL sockets. See the [“Database”](#) section for information about the connection between the Historian and the database. This process opens JMX port **4167** for monitoring. By default, the JMX port is not secured, but can be secured via SSL and username/password. See [“Monitored Components”](#) for information on securing this connection.

Database

The ALERTDEFS database stores alert threshold information and optionally alert persistence information. The Data Server connects to the ALERTDEFS database to query thresholds and also to set thresholds when the user interacts with the **Alert Administration** page in the user interface. The RTVHISTORY database stores cache data (if the Historian is enabled). The Historian connects to the RTVHISTORY database to insert cache history data and to perform data compaction. The Data Server connects to the RTVHISTORY database on startup to load initial history into the caches and also when the user interface asks for history data older than what is contained in the in-memory history caches.

By default, the Data Server and Historian will connect to the HSQLDB database that is included with RTView using an unsecured JDBC connection. See the Hsqldb documentation for information on configuring it for secure JDBC connections. Alternately, you can use your own database and secure the JDBC connection according to the documentation for that database.

Application Servers

The Solace PubSub+ Monitor comes with a Tomcat installation pre-configured with all of the necessary servlets. You can use this Tomcat or another Application Server. To deploy your servlets to your application server, go into the **projects/rtview-server** directory and run **update_wars.bat** or **update_wars.sh**. Copy all of the generated war files to the webapps directory in your application server.

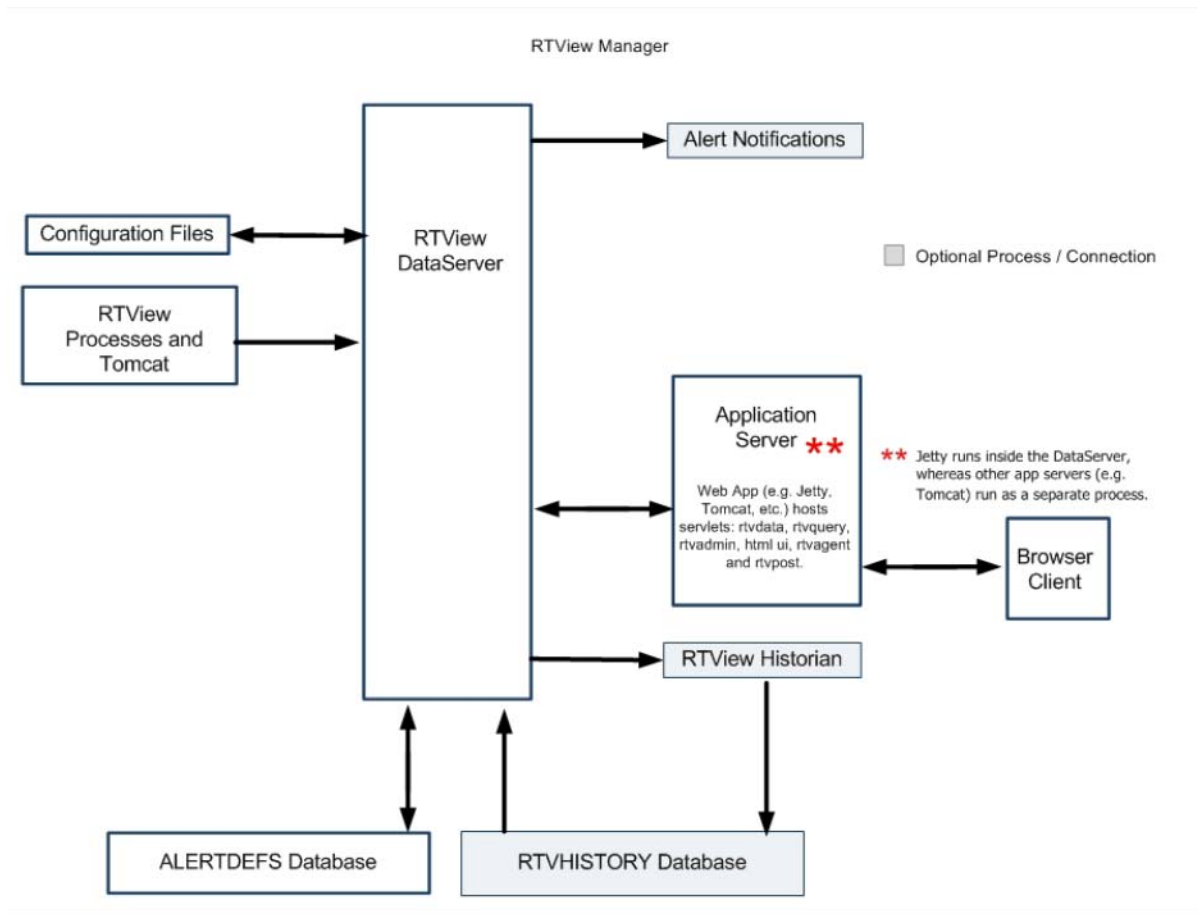
Tomcat and most other Application Servers can be configured for HTTPS. This will require you to provide a certificate for your domain. Follow the application server documentation for enable HTTPS. Additionally, Tomcat access filters can be configured to restrict access according to the remote client host or address. Tomcat also has a feature named LockOut Realm to protect against brute force login attacks. After 5 successive login attempts for a given username with invalid password, then all logins for that username are rejected for the next 5 minutes. The LockOut Realm parameters are configurable. See Tomcat documentation for more information.

If you are not using the Tomcat that came with Solace PubSub+ Monitor, you will need to add the following roles to your Application Server for use with the Configuration Application and HTML UI authentication. For Tomcat, users and roles are defined in **conf\tomcat-users.xml**:

- rtvadmin
- rtvuser
- rtvalertmgr

RTView Manager

Solace PubSub+ Monitor includes the RTView Manager which runs as a separate process to monitor RTView processes and Tomcat. The RTView Manager monitor is accessible in Tomcat at <http://host:port/rtview-manager> and also in Jetty that is running in the RTView Manager data server at <http://localhost:3070/rtview-manager>. The connections to the RTView processes and Tomcat are pre-configured. To modify other properties of the RTView Manager, go to the configuration application in Tomcat at <http://host:port/rtview-manager-rtvadmin> or in Jetty at <http://localhost:3070/rtview-manager-rtvadmin>.



The Tomcat application server included with this installation contains all of the RTView Manager servlets. You can use this Tomcat or another Application Server. To deploy your servlets to your application server, go into the **projects/rtview-manager** directory and run **update_wars.bat** or **update_wars.sh**. Copy all of the generated war files to the **webapps** directory in your application server.

The RTView Manager contains a Data Server, Historian, HTML UI, Configuration Application and Configuration files which can be secured as described in the sections above with the following exceptions:

- All configuration for RTView Manager is done in the **projects/rtview-manager** directory instead of the **rtview-server** directory.

The `rtvquery` servlet for the RTView Manager is not secured by default. To secure it, do the following:

- `cd` to `rtvapm` and `run rtvapm_init`
- `cd` to `projects\rtview-server`
- `run update_wars.bat` (or `.sh`) `-secure`
- deploy the generated war files to your Application Server

Note: Jetty does not support secured `rtvquery` servlets. You will need to use Tomcat or another Application Server.

Monitored Components

Monitored Components are the processes that the Data Server and Data Collector connect to in order to request metric data.

The Solace Data Server connects to cloud brokers via HTTP and non-cloud brokers via Solace API. See the Solace documentation for information on about securing your brokers.

To connect to a secured cloud broker, enter the HTTPS URL in the RTView Configuration Application Solace Connection dialog.

To connect to a secured non-cloud broker, turn on the **SSL Connection** toggle in the RTView Configuration Application Solace Connection dialog, then fill in the SSL credentials on the **SECURITY** tab of the RTView Configuration Application.

The RTView Manager Data Server connects to the Tomcat, Solace Data Server and Solace Historian via JMX. The processes that open JMX ports which can be configured to require a user name and password which the user enters in the RTView Configuration Application RTView Manager Connection dialog when defining the connection to that process. These processes can also be configured to require SSL. To connect to SSL secured JMX, fill in the SSL Credentials section of **SECURITY** tab in the RTView Configuration Application with the appropriate values for your SSL configuration.

NOTE: The Data Server, Data Collector, Historian and Display Server all open JMX ports for monitoring. By default, these JMX ports are unsecured, but they can be secured either by user name and password or by SSL.

Security Summary

Security options per RTView process are included in the section for each component above. This section provides a summary of security options for the entire deployment organized by priority.

This section contains:

- ["Secure Installation Location - High Priority"](#)
- ["Login and Servlet Authentication - High Priority"](#)
- ["Application Server Security - High Priority"](#)
- ["Secure Connections between RTView Processes - Medium-to-Low Priority*"](#)
- ["Secure Connections to Monitored Components - Medium-to-Low Priority*"](#)
- ["Secure Connections to Databases - Medium-to-Low Priority*"](#)

Secure Installation Location - High Priority

The RTView installation and Application Server should be run in a secure location to ensure displays and configuration files are secure and access-restricted.

Login and Servlet Authentication - High Priority

- **HTML UI** - By default, the HTML UI is configured with HTTP authentication which should be deployed on HTTPS since HTTP authentication does not encrypt user credentials. The HTML UI connects to the Data Server via the rtvquery servlet. The rtvquery servlet does not have authentication enabled by default. See the HTML UI section in this document for information on enabling authentication in the rtvquery servlet.
- **Configuration Application** - By default, the Configuration Application is configured with HTTP authentication which should use HTTPS since HTTP authentication does not encrypt user credentials.

Application Server Security - High Priority

It is highly recommended that you configure your Application Server to use HTTPS as described in the Application Server section of this document. The RTView servlets that support HTTP authentication do not encrypt user credentials.

It is highly recommended that you change the user credentials in your Application Server for the rtvadmin, rtvuser and rtvalertmgr roles since the default credentials are documented and publicly available.

Secure Connections between RTView Processes - Medium-to-Low Priority*

The Historian, Data Server, Data Collector, rtvquery servlet, rtvdata servlet, rtvadmin servlet and rtvagnt servlet all connect to the Data Server via socket which is unsecured by default. The Data Server supports secure socket connections (SSL) without certificates. Once the Data Server has been configured for secure sockets, the History and all rtvdata and rtvquery servlets will connect via secure socket.

Secure Connections to Monitored Components - Medium-to-Low Priority*

The Data Server uses component-specific APIs to connect to Monitored Components. Securing these connections is described here.

Secure Connections to Databases - Medium-to-Low Priority*

The Data Server and Historian both create database connections using JDBC. See the Database section in this document for information on securing JDBC connections to your database.

*If Secured Installation Location has been met, these are lower priority.

APPENDIX F 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 and pasted into an email.